

LARGE DEVIATIONS FOR WEIGHTED EMPIRICAL MEASURES ARISING IN IMPORTANCE SAMPLING

HENRIK HULT[†] AND PIERRE NYQUIST

ABSTRACT. Importance sampling is a popular method for efficient computation of various properties of a distribution such as probabilities, expectations, quantiles, etc. The output of an importance sampling algorithm can be represented as a weighted empirical measure, where the weights are given by the likelihood ratio between the original distribution and the sampling distribution. In this paper the efficiency of an importance sampling algorithm is studied by means of large deviations for the weighted empirical measure. The main result, which is stated as a Laplace principle for the weighted empirical measure arising in importance sampling, can be viewed as a weighted version of Sanov's theorem. The main theorem is applied to quantify the performance of an importance sampling algorithm over a collection of subsets of a given target set as well as quantile estimates. The proof of the main theorem relies on the weak convergence approach to large deviations developed by Dupuis and Ellis.

1. INTRODUCTION

Computational complexity is a central issue in the design of modern technology and systems. Cheaper and smaller devices enable us to collect and transmit huge amounts of data. The data can be utilized to make systems faster, safer, and more versatile. In order to use the data effectively we need fast and reliable techniques for analyzing the data and to perform advanced computational tasks. This paper contributes to the development of a new approach, based on the theory of large deviations for empirical measures, to the analysis of efficient computational methods within the context of stochastic simulation. In the present paper the emphasis is on efficiency and design for algorithms based on importance sampling.

Stochastic simulation is the collective term for simulating a physical system, involving random effects, on a computer. Computational methods based on stochastic simulation are fundamental in a wide range of applications, virtually all areas where probability is applied, including chemistry, computer science, finance, life sciences, networks, physics, power grids, reliability, solid mechanics, statistics, etc. The basic idea in stochastic simulation is to generate a population of particles that moves randomly according to the laws of the physical system. Each particle carries an individual weight, which may be updated during the simulation, and quantities of the underlying physical system are computed by averaging the particles' weights depending on their position. The standard example is Monte Carlo

2010 *Mathematics Subject Classification.* Primary 60F10, 65C05; secondary 60G57, 68U20.

Key words and phrases. Large deviations, empirical measures, importance sampling, Monte Carlo.

[†] Support from the Göran Gustafsson Foundation is gratefully acknowledged.

simulation where all the particles are independent and statistically identical and their weights are constant and equal.

Although the standard Monte Carlo procedure is widely used it is by no means universally applicable. One reason is that particles may wander off to irrelevant parts of the state space, leaving only a small fraction of relevant particles that contribute to the computational task at hand. Therefore standard Monte Carlo may require a huge number of particles to obtain a desired precision, resulting in a computational cost that is too high for practical purposes. A control mechanism is needed that forces the particles to move to the relevant part of the space, thereby increasing the importance of each particle and reducing the computational cost. The control mechanism may come in different form depending on the type of algorithm under consideration. In importance sampling, see e.g. [1], the control is the choice of sampling dynamics used to steer the particles towards the relevant part of the state space. In splitting algorithms [5] and interacting particle systems [2] the control mechanism come, roughly speaking, in the form of a birth/death mechanism, which controls that important particles give birth to new particles and irrelevant particles are killed.

The limited evidence provided by simply running numerical experiments has generated the need for a deeper theoretical understanding and analysis of the performance of stochastic simulation algorithms, see e.g. [1]. Much of the theoretical analysis on the efficiency of stochastic simulation algorithms in general, and importance sampling algorithms in particular, is based on analyzing the variance of the resulting estimators. The variance is the canonical measure of variability of unbiased estimators, but when an estimator is biased or skewed the variance can be misleading. This paper aims to complement the variance analysis by a detailed study of the rate function associated with a large deviation principle of the weighted empirical measure associated to the output of the algorithm. The main result is a Laplace principle for the weighted empirical measure resulting from a general importance sampling algorithm. The rate function associated to the Laplace principle can be used to identify what part of the design that is most likely to lead to computational errors and lead to a deeper understanding of how the design of an algorithm influences its performance.

Next follows a brief outline of our approach. For the sake of illustration, consider the problem of computing the probability of an event using importance sampling. In importance sampling the main design choice is the sampling dynamics to be used for generating the trajectories of the particles. The output of an importance sampling algorithm is a collection of particles at different locations with individual weights, represented as a weighted empirical measure. If the sampling dynamics are well chosen many particles will be located in or near the event and the variability of their weights will be small. As the number of particles increase the weighted empirical measure will look more and more like the true distribution on the event and the estimated probability of the event will converge to the true probability. But how fast? The theory of large deviations can be used to show that, under certain conditions, the error probability decays exponentially fast in the number of particles and the associated rate function tells us what the exponential rate is. The rate function depends on the design of the algorithm, in this case on the choice of sampling dynamics. The design choice then reduces to selecting sampling dynamics that maximize the exponential rate of decay of errors. Moreover, the rate function

will typically emphasize the features of a model that are most likely to contribute to estimation errors and discard those features that are of less importance. Thus, the rate function can potentially be used to identify key aspects of the sampling dynamics that will reduce the probability of errors. The key idea is that the properties of the rate function can be utilized to suggest new and improved algorithms.

To avoid confusion it should be pointed out that techniques from sample path large deviations have been studied thoroughly, to control the rate at which the variance decays, in the context of designing efficient rare-event simulation algorithms. Our objective is fundamentally different. We want to replace/complement the variance by the rate function of a large deviations principle as the number of particles increase. For our purposes the appropriate framework is large deviations for empirical measures, in the spirit of Sanov's theorem [4, Theorem 2.2.1], rather than sample path large deviations. The suggested approach is applicable to a wide range of simulation problems and is not intended exclusively for problems in rare-event simulation.

Our contributions can be summarized as follows. The main result in this paper is a Laplace principle, in the space of finite measures equipped with the τ -topology, for the weighted empirical measure arising in importance sampling. The result is expected in the sense that it can be guessed from Sanov's theorem and the contraction principle. Our proof of the general version, stated in Theorem 3.1, is based on the weak convergence approach to large deviations and follows, with some adaption, the proof of Sanov's theorem in [4]. Its relevance is mainly that it leads to a method for theoretical quantification of performance for importance sampling algorithms. The main theorem is applied to quantifying the performance of an importance sampling algorithm over a collection of subsets of a given target set as well as to quantile estimates. Furthermore, the result is potentially useful for theoretical comparison of the performance of algorithms of different character, based on, say, importance sampling and interacting particle systems, by comparing the associated rate functions.

The outline of the paper is as follows. In Section 2 an introduction to importance sampling is presented along with background on variance based efficiency analysis and large deviations for empirical measures. The main result, which is a Laplace principle for the weighted empirical measure of an importance sampling algorithm, is presented in Section 3. Applications of the main result to efficiency analysis and design of importance sampling algorithms are given in Section 4. Most of the examples are well studied elsewhere and are mainly intended to demonstrate the efficiency analysis using the rate function in contrast to the standard variance analysis. Section 5 contains the proof of the main theorem.

2. BACKGROUND

Let \mathcal{X} be a complete separable metric space equipped with its Borel σ -field $\mathcal{B}(\mathcal{X})$ and let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Unless otherwise stated subsets of \mathcal{X} under consideration are always assumed in $\mathcal{B}(\mathcal{X})$. Consider a random variable $X : \Omega \rightarrow \mathcal{X}$ with distribution F . Denote by \mathcal{M}_1 the space of probability measures on \mathcal{X} . The objective is to approximate F in a given region $A \in \mathcal{B}(\mathcal{X})$ or to compute $\Phi(F)$, where $\Phi : \mathcal{M}_1 \rightarrow \mathcal{R}$ is a given functional. Examples of functionals that may be of

interest are expectations,

$$\Phi_f(F) = \int f dF, \text{ for some } f : \mathcal{X} \rightarrow \mathcal{R},$$

and, when X is real valued, quantiles

$$\Phi_q(F) = F^{-1}(q) = \inf\{x : F((x, \infty)) \leq q\}, \quad q \in (0, 1),$$

and L -statistics,

$$\Phi(F) = \int_0^1 \phi(q) F^{-1}(q) dq.$$

Only in exceptional cases is explicit computation of such functionals possible. When explicit computations are not possible a viable alternative is simulation. The standard method of simulation is Monte Carlo, in which the empirical measure

$$\mathbf{F}_n = \frac{1}{n} \sum_{k=1}^n \delta_{X_k}$$

is constructed from an independent sample X_1, \dots, X_n from F . Here δ_x denotes a unit point mass at x . The quantity $\Phi(F)$ is estimated by the plug-in estimator $\Phi(\mathbf{F}_n)$. Roughly speaking, if \mathbf{F}_n is a good approximation of F in a region that largely determines $\Phi(F)$, then $\Phi(\mathbf{F}_n)$ is likely to provide a good estimate.

2.1. Large deviations analysis for quantifying the performance of Monte Carlo algorithms. Let us try to quantify how efficient the plug-in estimator is by means of large deviations for the associated empirical measure.

Obviously, the sample size n will affect the precision of the estimator $\Phi(\mathbf{F}_n)$. By the law of large numbers for empirical measures \mathbf{F}_n converges weakly to F with probability one, as $n \rightarrow \infty$, and an increased sample size n will thus improve the accuracy in this sense. In the case when $\Phi(\mathbf{F}_n)$ is an unbiased estimate of $\Phi(F)$ the sample size required to reach a desired precision is well captured by $\text{Var}(\Phi(F))$ and an analysis of the performance of an estimator can be done in terms of the variance. However, in the general case $\Phi(\mathbf{F}_n)$ can be biased and it might be that looking solely at the variance of the estimator may be insufficient.

An alternative way of quantifying the efficiency of the Monte Carlo estimator is through the theory of large deviations. To illustrate the point, let us consider the example of computing the expectation $F(f) = \int f dF$ for some $f : \mathcal{X} \mapsto \mathcal{R}$ by Monte Carlo simulation. An estimate of $F(f)$ is then given by

$$\mathbf{F}_n(f) = \frac{1}{n} \sum_{i=1}^n f(X_i), \tag{2.1}$$

where the X_i 's are independent with common distribution F .

Cramér's theorem states that if $\mathbb{E}[\exp\{\theta f(X)\}] < \infty$ for θ in a neighborhood of the origin, then $\mathbf{F}_n(f)$ satisfies the large deviation principle:

$$\begin{aligned} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\mathbf{F}_n(f) \in \overline{A}) &\leq -I(\overline{A}), \\ \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\mathbf{F}_n(f) \in A^\circ) &\geq -I(A^\circ), \end{aligned}$$

for all Borel sets A , where \overline{A} and A° denotes the closure and interior of A , respectively, $I(A) = \inf_{x \in A} I(x)$, and the rate function I is given by

$$I(x) = \sup_{\theta} \{\theta x - \kappa(\theta)\},$$

and $\kappa(\theta) = \log \mathbb{E}[\exp\{\theta f(X)\}]$ is the logarithm of the moment generating function [3, p. 26].

Suppose for the sake of illustration that, with probability at least $1 - \delta$, a relative precision ϵ is desired in the estimate. That is, the sample size n must be selected sufficiently high that

$$\mathbb{P}(|\mathbf{F}_n(f) - F(f)| \geq \epsilon F(f)) \leq \delta.$$

Cramér's theorem, with A_ϵ denoting the complement of the open ball of radius $\epsilon F(f)$ centered at $F(f)$, implies that

$$\limsup_n \frac{1}{n} \log \mathbb{P}(|\mathbf{F}_n(f) - F(f)| \geq \epsilon F(f)) = \limsup_n \frac{1}{n} \log \mathbb{P}(\mathbf{F}_n(f) \in A_\epsilon) \leq -I(A_\epsilon). \quad (2.2)$$

Then, at least approximately,

$$\mathbb{P}(\mathbf{F}_n(f) \in A_\epsilon) \lesssim e^{-nI(A_\epsilon)},$$

for large n , and the upper bound δ on the error probability corresponds to

$$n \gtrsim \frac{1}{I(A_\epsilon)} (-\log \delta). \quad (2.3)$$

Roughly speaking the sample size must be proportional to the reciprocal of the rate for the error probability in order to obtain the desired performance.

Let us consider a more general example where we are interested in approximating the true distribution F over a region $A \in \mathcal{B}(\mathcal{X})$. Again the empirical measure \mathbf{F}_n resulting from Monte Carlo simulation, restricted to A , provides a viable approximation. In this context a large deviation principle for the empirical measure can be applied to quantify the performance of the Monte Carlo algorithm. Sanov's theorem [4, Theorem 2.2.1] states that \mathbf{F}_n satisfies a large deviation principle on \mathcal{M}_1 with rate function given by the relative entropy $\mathcal{H}(\cdot | F)$, where

$$\mathcal{H}(G | F) = \int \left(\log \frac{dG}{dF} \right) dG, \quad G \in \mathcal{M}_1.$$

More precisely,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\mathbf{F}_n \in \overline{A}) &\leq -I(\overline{A}), \\ \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\mathbf{F}_n \in A^\circ) &\geq -I(A^\circ), \end{aligned}$$

where $A \in \mathcal{B}(\mathcal{M}_1)$ and $I(G) = \mathcal{H}(G | F)$.

Taking $A = U_G$ to be a small neighborhood of G one can interpret, when n is large, $\exp\{-nI(G)\}$ as an approximation for the probability that the empirical measure looks like a typical sample from G . If it is undesirable that the empirical measure looks like a typical sample from a specific distribution G , say the objective is to have

$$\mathbb{P}(\mathbf{F}_n \in U_G) \leq \delta,$$

then, by the same reasoning as above, the sample size must be selected sufficiently large that

$$n \gtrsim \frac{1}{I(G)}(-\log \delta). \quad (2.4)$$

Notice that taking $A = \{G \in \mathcal{M}_1 : |G(f) - F(f)| \geq \epsilon F(f)\}$ in Sanov's theorem one can recover (2.3).

In the context of a general functional Φ Sanov's theorem may also be applied. Either one would try to protect against a certain undesirable shape G of the empirical measure by selecting n as in (2.4) or by looking at error probabilities corresponding to sets of the form

$$A_\epsilon = \{G \in \mathcal{M}_1 : |\Phi(G) - \Phi(F)| \geq \epsilon \Phi(F)\}.$$

It may be pointed out that when considering a set such as A_ϵ the rate is given by

$$I(A_\epsilon) = \inf_{G \in A_\epsilon} I(G),$$

so if the infimum is attained at G_ϵ , then limiting the probability of A_ϵ corresponds precisely to protecting against G_ϵ .

2.2. Importance sampling. Importance sampling is a popular method to improve the accuracy of Monte Carlo simulation. The basic idea is to draw the samples from a sampling distribution that is more likely to generate samples from the desired region. Suppose that the goal is to evaluate $\Phi(F)$ for some functional Φ . For simplicity, start by considering the case $\Phi(F) = F(f) = \int f(x)F(dx)$ for an F -integrable, non-negative function $f : \mathcal{X} \rightarrow \mathcal{R}$. Let \tilde{F} be the chosen sampling distribution. For \tilde{F} to be a feasible sampling distribution it must hold that $F \ll \tilde{F}$ on the support of f . Then the Radon-Nikodym derivative $dF/d\tilde{F}$ exists on $\{f > 0\}$ and it is possible to define the weight function

$$w(x) = \frac{dF}{d\tilde{F}}(x)I\{f(x) > 0\}. \quad (2.5)$$

Let $\tilde{X}_1, \dots, \tilde{X}_n$ be an independent sample from \tilde{F} . The weighted empirical measure corresponding to the importance sampling algorithm is

$$\tilde{\mathbf{F}}_n^w = \frac{1}{n} \sum_{k=1}^n w(\tilde{X}_k) \delta_{\tilde{X}_k}.$$

Note that in contrast to standard Monte Carlo $\tilde{\mathbf{F}}_n^w$ is typically not a probability measure. The importance sampling estimator of $F(f)$ is the plug-in estimator

$$\tilde{\mathbf{F}}_n^w(f) = \frac{1}{n} \sum_{i=1}^n w(\tilde{X}_i) f(\tilde{X}_i). \quad (2.6)$$

Let $\tilde{\mathbb{P}}$ and $\tilde{\mathbb{E}}$ denote the probability and expectation when $\tilde{X}_1, \tilde{X}_2, \dots$ are sampled from the sampling distribution \tilde{F} . If $\tilde{\mathbb{E}}[\exp\{\theta w(\tilde{X})f(\tilde{X})\}] < \infty$ for θ in a neighborhood of the origin, then Cramér's theorem implies that $\tilde{\mathbf{F}}_n^w(f)$ satisfies a large deviation principle with rate function

$$I^w(x) = \sup_{\theta} \{\theta x - \kappa^w(\theta)\},$$

$$\kappa^w(\theta) = \log \tilde{\mathbb{E}}[\exp\{\theta w(\tilde{X})f(\tilde{X})\}].$$

Suppose, as for the Monte Carlo illustration that, with probability at least $1 - \delta$, a relative precision ϵ is desired in the estimate. That is, the sample size n must be selected sufficiently high that

$$\widetilde{\mathbb{P}}(|\tilde{\mathbf{F}}_n^w(f) - F(f)| \geq \epsilon F(f)) \leq \delta.$$

Just as before Cramér's theorem, with A_ϵ denoting the complement of the open ball of radius $\epsilon F(f)$ centered at $F(f)$, implies that the sample size must satisfy

$$n \gtrsim \frac{1}{I^w(A_\epsilon)}(-\log \delta).$$

The choice of the sampling distribution enters in the rate function through the weights $w = (dF/d\tilde{F})I\{f > 0\}$. The improvement over standard Monte Carlo can be quantified by comparing the rate function I corresponding to Monte Carlo and the rate function I^w corresponding to importance sampling. Furthermore, a good choice of the sampling distribution in the importance sampling algorithm is one that maximizes the rate $I^w(A_\epsilon)$.

To extend the analysis to more general functionals it is desirable to have an analogue of Sanov's theorem for the weighted empirical measures $\tilde{\mathbf{F}}_n^w$. In contrast to Monte Carlo one cannot expect that the weighted empirical measure $\tilde{\mathbf{F}}_n^w$ is a good approximation to F everywhere. Rather the sampling distribution is selected to obtain a good precision in the important part of the space. If the objective is to compute $\Phi(F)$ for some functional Φ , then it suffices that $\tilde{\mathbf{F}}_n^w$ approximates F in the region that largely determines $\Phi(F)$. For this purpose a non-negative measurable function $f : \mathcal{X} \rightarrow \mathcal{R}$ is introduced, called the *importance function*.

The rough statement that $\tilde{\mathbf{F}}_n^w$ is close to F in the important region is made precise by saying that the measure $\tilde{\mathbf{F}}_n^{wf}$ is close to F^f in the space $\mathcal{M} = \mathcal{M}(\mathcal{X})$ of finite measures, where

$$\tilde{\mathbf{F}}_n^{wf} = \frac{1}{n} \sum_{k=1}^n w(X_k) f(X_k) \delta_{X_k},$$

and F^f is the finite measure given by

$$F^f(g) = \int g(x) f(x) F(dx),$$

for each bounded measurable $g : \mathcal{X} \rightarrow \mathcal{R}$. To establish an analogue of Sanov's theorem for the weighted empirical measures is the main objective in this paper.

3. A LAPLACE PRINCIPLE FOR WEIGHTED EMPIRICAL MEASURES

In this section the main result of this paper is stated. It is an extension of Sanov's theorem to the weighted empirical measures arising in importance sampling, stated as a Laplace principle.

A sequence of random variables U_n taking values in a topological space \mathcal{U} is said to satisfy a *Laplace principle* on \mathcal{U} with rate function I if, for all bounded, continuous functions $h : \mathcal{U} \rightarrow \mathcal{R}$, it satisfies

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E}[e^{-nh(U_n)}] = - \inf_{u \in \mathcal{U}} \{h(u) + I(u)\}. \quad (3.1)$$

When \mathcal{U} is a Polish space, the Laplace principle is equivalent to the sequence satisfying a large deviation principle with the same rate function [4, Theorems 1.2.1, 1.2.3]. In the case of a general topological space, the relationship between the large

deviation principle and the Laplace principle is given by Varadhan's lemma and Bryc's inverse, see [3] and the references therein.

Suppose that F and \tilde{F} are two given distributions on a complete, separable metric space \mathcal{X} . The space \mathcal{M} of finite measures on \mathcal{X} is equipped with the τ -topology; $\nu_n \xrightarrow{\tau} \nu$ if $\nu_n(g) \rightarrow \nu(g)$ for all bounded measurable $g : \mathcal{X} \rightarrow \mathcal{R}$. To avoid the subtle measurability issues of the τ -topology, see [4] pp. 333-334, we will have to work with $\mathcal{F}_{\mathcal{M}}$ – the smallest σ -field on \mathcal{M} with respect to which the function mappings $\nu \mapsto \int g d\nu$ are measurable.

Let f be an importance function. That is, f is a non-negative F -integrable function characterizing the importance of different regions of \mathcal{X} . It is assumed that $F \ll \tilde{F}$ on the support of f and the weight function w is defined as in (2.5). Let \mathcal{M}_1 , the space of probability measures, be equipped with the τ -topology as well and introduce the set $\Gamma = \{G \in \mathcal{M}_1 : G(wf) < \infty\}$. Define the mapping Ψ from the subset $\Gamma \subset \mathcal{M}_1$ to \mathcal{M} as the mapping for which $\Psi(G; \cdot)$ is the finite measure given by

$$\Psi(G; g) = \int g(x) f(x) w(x) G(dx), \quad (3.2)$$

for each bounded measurable $g : \mathcal{X} \rightarrow \mathcal{R}$. A key observation is that $\tilde{\mathbf{F}}_n^{wf}(g) = \tilde{\mathbf{F}}_n(wfg)$, where

$$\tilde{\mathbf{F}}_n = \frac{1}{n} \sum_{k=1}^n \delta_{\tilde{X}_k},$$

is the empirical measure obtained by sampling from \tilde{F} , and therefore $\tilde{\mathbf{F}}_n^{wf} = \Psi(\tilde{\mathbf{F}}_n; \cdot)$. Note also that $\tilde{\mathbf{F}}_n$ belongs to Γ with probability 1.

Let $\Delta \subset \mathcal{M}_1$ be the set

$$\Delta = \{G \in \mathcal{M}_1 : \mathcal{H}(G | \tilde{F}) < \infty\} \quad (3.3)$$

We are now ready to introduce the rate function. Let $I : \mathcal{M} \mapsto [0, \infty]$ be the function defined by

$$I(\nu) = \inf\{\mathcal{H}(G | \tilde{F}) : G \in \Gamma \cap \Delta, \Psi(G) = \nu\}, \quad (3.4)$$

when such G exist and $I(\nu) = \infty$ otherwise. Proposition 5.10 below states that I has sequentially compact level sets. Our main result is the following.

Theorem 3.1 (Laplace principle for weighted empirical measures). *Let F and \tilde{F} be given as above and let f be an importance function. Suppose that*

- (i) *there exists a function $U : \mathcal{X} \rightarrow [0, \infty]$ such that $\int e^{U(x)} \tilde{F}(dx) < \infty$ and U has relatively compact level sets,*
- (ii) *$\int e^{\alpha w(x) f(x)} d\tilde{F}(x) < \infty$ for all $\alpha > 0$.*

Then the sequence $\{\tilde{\mathbf{F}}_n^{wf}\}$ of weighted empirical measures satisfies the following Laplace principle on \mathcal{M} equipped with the τ -topology. For all bounded continuous functions $h : \mathcal{M} \rightarrow \mathcal{R}$, measurable on $(\mathcal{M}, \mathcal{F}_{\mathcal{M}})$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \tilde{\mathbb{E}}[e^{-nh(\tilde{\mathbf{F}}_n^{wf})}] = - \inf_{\nu \in \mathcal{M}} \{h(\nu) + I(\nu)\}, \quad (3.5)$$

with the rate function I in (3.4).

The proof is given in Section 5.

Remark 3.2. Condition (i) is a version of Condition 8.2.2 in [4], adapted to the case of independent and identically distributed variables. In the case of real-valued random variables it is a very mild assumption. Take, for example,

$$U(x) = \max\{\alpha \log(|x|), 0\},$$

for some $\alpha > 0$. Then U has relatively compact level sets and the condition of a finite expectation of $e^{U(X)}$ is in this case weaker than

$$\mathbb{E}[|X|^\alpha] < \infty,$$

for some $\alpha > 0$.

Remark 3.3. The right-hand side of (3.5) can be written

$$\inf_{\nu \in \mathcal{M}} \{h(\nu) + I(\nu)\} = \inf_{G \in \Gamma \cap \Delta} \{h(\Psi(G)) + \mathcal{H}(G|\tilde{F})\}.$$

Although the expression on the left is the standard way to express the limit of the Laplace principle, the expression on the right better suits the weak convergence approach to large deviations that will be adopted throughout the proof.

Consider the special case when the function wf is bounded. Then, $\Psi : \Gamma \rightarrow \mathcal{M}$ is continuous when both spaces are equipped with the τ topology and Theorem 3.1 follows essentially from a standard application of the contraction principle, see e.g. [4, Theorem 1.3.2]. The main difficulty is to show that the Laplace principle holds also in the general case.

4. APPLICATIONS IN PERFORMANCE ANALYSIS

In this section the Laplace principle of Theorem 3.1 is applied to characterize the performance of an importance sampling algorithm. In the first part we outline a method for analyzing the performance over a collection of subsets of a target region A . If the sampling distribution is designed for a target set A , then one can expect to have good performance for subsets C of A that are not too small relative to A . A few rather elementary examples illustrate the performance analysis based on Theorem 3.1. We also discuss briefly the rare-event limit when the target set has small probability. The section ends with a brief discussion on performance analysis for importance sampling algorithms designed for computing the quantile of a distribution.

4.1. Performance over a collection of subsets. In this section we are interested in the performance of importance sampling algorithms over a region $A \subset \mathcal{X}$, reflected in the importance function $f(x) = I\{x \in A\}$. The ideal is that the weighted empirical measure $\tilde{\mathbf{F}}_n^{wf}$ is close to F on all measurable subsets of A . For large n we can imagine that the weighted empirical measure looks like typical sample from a measure ν , which is absolutely continuous with respect to F . The performance of the importance sampling algorithm is good if it is likely that $\tilde{\mathbf{F}}_n^{wf}$ looks like a typical sample from some ν belonging to a set of measures for which the likelihood ratio $d\nu/dF$ is close to 1 on A . For given $\epsilon > 0$ and $\delta > 0$ (where $\delta = \delta' F(A)$ for some $\delta' > 0$ is a reasonable choice), consider the sets

$$\begin{aligned} A_{\epsilon, \delta} &= \{\nu \in \mathcal{M} : |d\nu/dF(x) - 1| \geq \epsilon \text{ for } x \in \text{some } C \subset A \text{ with } F(C) \geq \delta\}, \\ A_{\epsilon, \delta}^+ &= \{\nu \in \mathcal{M} : d\nu/dF(x) \geq 1 + \epsilon \text{ for } x \in \text{some } C \subset A \text{ with } F(C) \geq \delta\}, \end{aligned} \quad (4.1)$$

$$A_{\epsilon, \delta}^- = \{\nu \in \mathcal{M} : d\nu/dF(x) \leq 1 - \epsilon \text{ for } x \in \text{some } C \subset A \text{ with } F(C) \geq \delta\}. \quad (4.2)$$

The rate $I(A_{\epsilon,\delta})$, with I as in Theorem 3.1, can be used to evaluate the performance of the importance sampling algorithm. The interpretation is that $e^{-nI(A_{\epsilon,\delta})}$ is roughly the probability that $\tilde{\mathbf{F}}_n^{wf}$ provides an approximation of F with relative error greater than ϵ for all subsets $C \subset A$ that are not too small in the sense that $F(C) \geq \delta$. The sets $A_{\epsilon,\delta}^+$ and $A_{\epsilon,\delta}^-$ have similar interpretations for overestimation and underestimation, respectively. The sets $A_{\epsilon,\delta}^+$ and $A_{\epsilon,\delta}^-$ are somewhat easier to analyze than $A_{\epsilon,\delta}$, so for the sake of illustration they will be studied throughout the rest of this section. In the examples that follows we will, to keep it short, work exclusively with $A_{\epsilon,\delta}^+$.

If G is a probability measure such that $\nu = \Psi(G)$, then $d\nu/dF = dG/d\tilde{F}$ on A and

$$I(A_{\epsilon,\delta}^+) = \inf \left\{ \mathcal{H}(G \mid \tilde{F}) : \frac{dG}{d\tilde{F}}(x) \geq 1 + \epsilon \text{ for } x \in \text{some } C \subset A \text{ with } F(C) \geq \delta \right\}.$$

Let us compute $I(A_{\epsilon,\delta}^+)$. To start off, consider a fixed set $C \subset A$.

Lemma 4.1. *Given $C \subset A$ it holds that*

$$\begin{aligned} & \inf \left\{ \mathcal{H}(G \mid \tilde{F}) : \frac{dG}{d\tilde{F}}(x) \geq 1 + \epsilon \text{ for } x \in C \right\} \\ &= (1 + \epsilon)\tilde{F}(C) \log(1 + \epsilon) + (1 - (1 + \epsilon)\tilde{F}(C)) \log \left(\frac{1 - (1 + \epsilon)\tilde{F}(C)}{1 - \tilde{F}(C)} \right), \end{aligned}$$

where the infimum is attained for the probability measure G^* with

$$\begin{aligned} \frac{dG^*}{d\tilde{F}}(x) &= 1 + \epsilon, \text{ for } x \in C, \\ \frac{dG^*}{d\tilde{F}}(x) &= \frac{1 - (1 + \epsilon)\tilde{F}(C)}{1 - \tilde{F}(C)}, \text{ for } x \in C^c. \end{aligned}$$

Proof. For any probability measure G with $\frac{dG}{d\tilde{F}}(x) \geq 1 + \epsilon$ on C , convexity of the function $\varphi(s) = s \log s$ and Jensen's inequality implies that

$$\begin{aligned} \mathcal{H}(G \mid \tilde{F}) &= \tilde{F}(C) \int_C \varphi\left(\frac{dG}{d\tilde{F}}(x)\right) \frac{\tilde{F}(dx)}{\tilde{F}(C)} + \tilde{F}(C^c) \int_{C^c} \varphi\left(\frac{dG}{d\tilde{F}}(x)\right) \frac{\tilde{F}(dx)}{\tilde{F}(C^c)} \\ &\geq \tilde{F}(C) \left(\frac{G(C)}{\tilde{F}(C)} \log \frac{G(C)}{\tilde{F}(C)} \right) + \tilde{F}(C^c) \left(\frac{G(C^c)}{\tilde{F}(C^c)} \log \frac{G(C^c)}{\tilde{F}(C^c)} \right) \\ &= G(C) (\log G(C) - \log \tilde{F}(C)) + (1 - G(C)) [\log(1 - G(C)) - \log(1 - \tilde{F}(C))]. \end{aligned}$$

The last expression is convex as a function of $G(C)$ and is minimized at $G(C) = (1 + \epsilon)\tilde{F}(C)$. We conclude that the lower bound

$$\begin{aligned} & \inf \left\{ \mathcal{H}(G \mid \tilde{F}) : \frac{dG}{d\tilde{F}}(x) \geq 1 + \epsilon \text{ for } x \in C \right\} \\ &\geq (1 + \epsilon)\tilde{F}(C) \log(1 + \epsilon) + (1 - (1 + \epsilon)\tilde{F}(C)) \log \left(\frac{1 - (1 + \epsilon)\tilde{F}(C)}{1 - \tilde{F}(C)} \right), \end{aligned}$$

holds. It is straightforward to check that the lower bound is attained by G^* . This completes the proof. \square

Denote by J_+ the set function given by

$$J_+(C) = \inf \left\{ \mathcal{H}(G | \tilde{F}) : \frac{dG}{d\tilde{F}}(x) \geq 1 + \epsilon \text{ for } x \in C \right\}.$$

The rate $I(A_{\epsilon,\delta}^+)$ can be computed by minimizing the function J_+ over the feasible sets.

Lemma 4.2. $I(A_{\epsilon,\delta}^+) = \inf \{ J_+(C) : C \subset A, F(C) \geq \delta \}$.

Proof. For any $C^* \subset A$ such that $F(C^*) \geq \delta$ we have

$$\begin{aligned} I(A_{\epsilon,\delta}^+) &= \inf \left\{ \mathcal{H}(G | \tilde{F}) : \frac{dG}{d\tilde{F}}(x) \geq 1 + \epsilon \text{ for } x \in \text{some } C \subset A \text{ with } F(C) \geq \delta \right\} \\ &\leq \inf \left\{ \mathcal{H}(G | \tilde{F}) : \frac{dG}{d\tilde{F}}(x) \geq 1 + \epsilon \text{ for } x \in C^* \right\} \\ &= J_+(C^*). \end{aligned}$$

Taking infimum over feasible sets C^* leads to

$$I(A_{\epsilon,\delta}^+) \leq \inf \{ J_+(C^*) : C^* \subset A, F(C^*) \geq \delta \}.$$

It remains to show the reverse inequality. For every $\eta > 0$ there exists a probability measure $G^* \in \Delta \cap \Gamma$ and a corresponding set $C^* \subset A$, with $F(C^*) \geq \delta$, such that

$$I(A_{\epsilon,\delta}^+) + \eta \geq \mathcal{H}(G^* | \tilde{F}) \geq J_+(C^*) \geq \inf \{ J_+(C) : C \subset A \text{ and } F(C) \geq \delta \}.$$

Since $\eta > 0$ is arbitrary the proof is complete. \square

The next result characterizes the minimizing set C in Lemma 4.2 in terms of the likelihood ratio.

Lemma 4.3. For any $t \geq 0$, let

$$C_t = \left\{ x \in A : \frac{dF}{d\tilde{F}}(x) \geq t \right\},$$

and $\delta > 0$. If there exists \tilde{t}_δ such that $F(C_{\tilde{t}_\delta}) = \delta$, then the infimum in Lemma 4.2 is attained by $C_{\tilde{t}_\delta}$. That is,

$$I(A_{\epsilon,\delta}^+) = \inf \{ J_+(C) : C \subset A, F(C) \geq \delta \} = J_+(C_{\tilde{t}_\delta}).$$

In general,

$$\sup \{ J_+(C_t) : t \geq 0, F(C_t) \leq \delta \} \leq I(A_{\epsilon,\delta}^+) \leq \inf \{ J_+(C_t) : t \geq 0, F(C_t) \geq \delta \}.$$

Proof. The expression for $\mathcal{H}(G^* | \tilde{F})$ in Lemma 4.1 is increasing in $\tilde{F}(C)$. Thus, minimizing $J_+(C)$ corresponds to taking infimum of $\tilde{F}(C)$. Consider the problem

$$\inf \{ \tilde{F}(C) : C \subset A, F(C) \geq \delta \}.$$

Let \mathcal{F}_A be the collection of all measurable functions $\alpha : A \mapsto [0, 1]$. Since indicator functions of subsets of A are included in \mathcal{F}_A ,

$$\inf \{ \tilde{F}(C) : C \subset A, F(C) \geq \delta \} \geq \inf \left\{ \int_A \alpha(x) \tilde{F}(dx) : \alpha \in \mathcal{F}_A, \int_A \alpha(x) F(dx) \geq \delta \right\},$$

with equality if the infimum on the right-hand-side is attained for an indicator function.

Suppose that there exists \tilde{t}_δ with $F(C_{\tilde{t}_\delta}) = \delta$. We claim that $\alpha^*(x) = I\{x \in C_{\tilde{t}_\delta}\}$ is the solution to the problem on the right in the last display. To see this, take an arbitrary $\alpha \in \mathcal{F}_A$ such that $\int_A \alpha(x)F(dx) \geq \delta$.

$$\begin{aligned}
\int_A (\alpha(x) - \alpha^*(x))\tilde{F}(dx) &= \int_A (\alpha(x) - \alpha^*(x))\frac{d\tilde{F}}{dF}(x)F(dx) \\
&= \int_{C_{\tilde{t}_\delta}} (\alpha(x) - 1)\frac{d\tilde{F}}{dF}(x)F(dx) + \int_{A \setminus C_{\tilde{t}_\delta}} \alpha(x)\frac{d\tilde{F}}{dF}(x)F(dx) \\
&\geq \int_{C_{\tilde{t}_\delta}} (\alpha(x) - 1)\frac{1}{\tilde{t}_\delta}F(dx) + \int_{A \setminus C_{\tilde{t}_\delta}} \alpha(x)\frac{1}{\tilde{t}_\delta}F(dx) \\
&= \frac{1}{\tilde{t}_\delta} \int_A \alpha(x)F(dx) - \frac{1}{\tilde{t}_\delta} \int_{C_{\tilde{t}_\delta}} F(dx) \\
&\geq 0,
\end{aligned}$$

by the requirements on α and \tilde{t}_δ .

In the general case, it is obvious that

$$\inf\{J_+(C) : C \subset A, F(C) \geq \delta\} \leq \inf\{J_+(C_t) : t \text{ such that } F(C_t) \geq \delta\}.$$

Moreover, for any $t \geq 0$ such that $F(C_t) \leq \delta$ the same arguments as above yields

$$\int_A (\alpha(x) - I\{x \in C_t\})\tilde{F}(dx) \geq \frac{1}{t} \int_A \alpha(x)F(dx) - \frac{1}{t}F(C_t) \geq 0,$$

for all $\alpha \in \mathcal{F}_A$ with $\int_A \alpha(x)F(dx) \geq \delta$. We conclude that

$$\inf\left\{\int_A \alpha(x)\tilde{F}(dx) : \alpha \in \mathcal{F}_A, \int_A \alpha(x)F(dx) \geq \delta\right\} \geq \tilde{F}(C_t),$$

and as a consequence that

$$\inf\{J_+(C) : C \subset A, F(C) \geq \delta\} \geq J_+(C_t).$$

The proof is completed by taking the supremum over t such that $F(C_t) \leq \delta$. \square

Denote by γ_ϵ^+ the function

$$\gamma_\epsilon^+(s) = (1 + \epsilon)s \log(1 + \epsilon) + (1 - (1 + \epsilon)s) \log\left(\frac{1 - (1 + \epsilon)s}{1 - s}\right).$$

so that the expression in Lemma 4.1 coincides with $\gamma_\epsilon^+(\tilde{F}(C))$. Then,

$$I(A_{\epsilon, \delta}^+) = \gamma_\epsilon^+(\tilde{F}(C_{\tilde{t}_\delta})),$$

where \tilde{t}_δ is such that $\tilde{F}(C_{\tilde{t}_\delta}) = \delta$. Observe that γ_ϵ^+ is an increasing function. Therefore, a good choice of sampling distribution is one that makes $\tilde{F}(C_{\tilde{t}_\delta})$ large.

Next, consider the set $A_{\epsilon, \delta}^-$ in (4.2). The following results are obtained completely analogously to the case $A_{\epsilon, \delta}^+$. J_- is the direct analogue of the set function J_+ .

Lemma 4.4. *Given $C \subset A$ it holds that*

$$\begin{aligned}
&\inf\left\{\mathcal{H}(G \mid \tilde{F}) : \frac{dG}{d\tilde{F}}(x) \leq 1 - \epsilon \text{ for } x \in C\right\} \\
&= (1 - \epsilon)\tilde{F}(C) \log(1 - \epsilon) + (1 - (1 - \epsilon)\tilde{F}(C)) \log\left(\frac{1 - (1 - \epsilon)\tilde{F}(C)}{1 - \tilde{F}(C)}\right),
\end{aligned}$$

where the infimum is attained for the probability measure G_* with

$$\begin{aligned}\frac{dG_*}{d\tilde{F}}(x) &= 1 - \epsilon, \text{ for } x \in C, \\ \frac{dG_*}{d\tilde{F}}(x) &= \frac{1 - (1 - \epsilon)\tilde{F}(C)}{1 - \tilde{F}(C)}, \text{ for } x \in C^c.\end{aligned}$$

Lemma 4.5.

$$I(A_{\epsilon,\delta}^-) = \inf\{J_-(C) : C \subset A, F(C) \geq \delta\}.$$

Let γ_ϵ^- be the function

$$\gamma_\epsilon^-(s) = (1 - \epsilon)s \log(1 - \epsilon) + (1 - (1 - \epsilon)s) \log\left(\frac{1 - (1 - \epsilon)s}{1 - s}\right).$$

This is an increasing function in s and thus the optimal C is given precisely by the $C_{\tilde{t}_\delta}$ in Lemma 4.3. Thus, $I(A_{\epsilon,\delta}^-) = \gamma_\epsilon^-(\tilde{F}(C_{\tilde{t}_\delta}))$, whenever there exists \tilde{t}_δ such that $F(C_{\tilde{t}_\delta}) = \delta$.

Example 4.6. Consider a standard Monte Carlo algorithm (i.e. $\tilde{F} = F$) and let A be a set with $F(A) = p$. Put

$$A_{\epsilon,\delta}^+ = \left\{ G \in \mathcal{M}_1 : \frac{dG}{dF}(x) \geq 1 + \epsilon \text{ for } x \in \text{some } C \subset A, F(C) \geq \delta \right\}.$$

The rate function, given by Sanov's theorem, is the relative entropy: $I^{MC}(G) = \mathcal{H}(G \mid F)$. The rate can be computed, just as in the general importance sampling case, as

$$I^{MC}(A_{\epsilon,\delta}^+) = \inf\{J_+^{MC}(C) : C \subset A, F(C) \geq \delta\},$$

with

$$J_+^{MC}(C) = \inf\{\mathcal{H}(G \mid F) : \frac{dG}{dF}(x) \geq 1 + \epsilon, x \in C\} = \gamma_\epsilon^+(F(C)).$$

Suppose there exists a set C such that $F(C) = \delta$. Then, since γ_ϵ^+ is increasing, we conclude that

$$I^{MC}(A_{\epsilon,\delta}^+) = \gamma_\epsilon^+(\delta),$$

and, by the reasoning leading to (2.3), that the number of samples needed in order to obtain a specific error probability is proportional to the reciprocal of the rate. With I^{IS} denoting the rate function of Theorem 3.1, under the assumption that in Lemma 4.3 such a \tilde{t}_δ exists, an importance sampling algorithm with sampling distribution \tilde{F} has rate

$$I^{IS}(A_{\epsilon,\delta}^+) = \gamma_\epsilon^+(\tilde{F}(C_{\tilde{t}_\delta})).$$

If the cost for generating one sample from \tilde{F} is c times the cost for generating one sample from F , then the reduction in computational cost is roughly

$$c \frac{n^{IS}}{n^{MC}} \approx c \frac{I^{MC}(A_{\epsilon,\delta}^+)}{I^{IS}(A_{\epsilon,\delta}^+)} = c \frac{\gamma_\epsilon^+(\delta)}{\gamma_\epsilon^+(\tilde{F}(C_{\tilde{t}_\delta}))}.$$

Example 4.7 (Light-tailed random walk). Let F denote the distribution of a normalized light-tailed random walk with m steps. That is, F is the distribution of $S_m/m = (X_1 + \dots + X_m)/m$ where X_1, \dots, X_m are independent and identically distributed with finite moment generating function. Take $A = (a, \infty)$ to be the set of interest, i.e. $F(A) = \mathbb{P}(S_m \geq ma)$ and take the sampling distribution \tilde{F}_θ as an exponential change of measure:

$$\frac{dF}{d\tilde{F}_\theta}(x_1, \dots, x_m) = e^{m(\kappa(\theta) - \theta \frac{1}{m} \sum_{i=1}^m x_i)}.$$

Here κ is the log moment generating function of X_1 . Consider the set $A_{\epsilon, \delta}^+$ in (4.1). From Lemma 4.3 the rate $I(A_{\epsilon, \delta}^+)$ is given by $\gamma_\epsilon^+(\tilde{F}_\theta(C_{\tilde{t}_\delta}))$ where

$$C_{\tilde{t}_\delta} = \left\{ x \in A : \frac{1}{m} \sum_{i=1}^m x_i \leq (\kappa(\theta) - (1/m) \log \tilde{t}_\delta) / \theta \right\}.$$

Note that, by the choice of \tilde{t}_δ ,

$$\begin{aligned} \tilde{F}_\theta(C_{\tilde{t}_\delta}) &= \mathbb{E}[e^{\theta S_m - m\kappa(\theta)} I\{a \leq S_m/m \leq (\kappa(\theta) - (1/m) \log \tilde{t}_\delta) / \theta\}] \\ &\geq e^{m(\theta a - \kappa(\theta))} \mathbb{E}[I\{a \leq S_m/m \leq (\kappa(\theta) - (1/m) \log \tilde{t}_\delta) / \theta\}] \\ &= e^{m(\theta a - \kappa(\theta))} F(C_{\tilde{t}_\delta}) \\ &= e^{m(\theta a - \kappa(\theta))} \delta. \end{aligned}$$

If the cost for each replication of the importance sampling algorithm is c times the cost for each replication of the standard Monte Carlo algorithm, then we conclude that the reduction in computational cost is given by

$$c \frac{\gamma_\epsilon^+(\delta)}{\gamma_\epsilon^+(\tilde{F}(C_{\tilde{t}_\delta}))} \leq c \frac{\gamma_\epsilon^+(\delta)}{\gamma_\epsilon^+(e^{m(\theta a - \kappa(\theta))} \delta)}.$$

A good choice of θ is the maximizer to $\theta a - \kappa(\theta)$, which is given by θ_a such that $\kappa'(\theta_a) = a$. In addition to suggesting the well known exponential change of measure with parameter θ_a our large deviations analysis also provides a useful upper bound on the reduction in computational cost.

4.2. Applications to rare-event simulation. The efficiency analysis of the importance sampling algorithms presented so far is not targeted specifically to capture the performance of rare-event simulation algorithms. In this section we illustrate how a rare-event analysis can be performed, based on Theorem 3.1. The elementary examples presented in this section demonstrate that, based on Theorem 3.1, one can obtain similar results on rare-event efficiency as in the standard case where the efficiency analysis is based on the variance.

We begin by analyzing standard Monte Carlo and the importance sampling algorithm for the light-tailed random walk with an emphasis on rare events.

Example 4.8 (Rare-event analysis for standard Monte Carlo). Consider computing a probability $F(A) = p$ by standard Monte Carlo, as in Example 4.6. Take $\delta = \delta' p$ for some $\delta' \in (0, 1)$. The performance of the algorithm can be captured by the rate

$I^{MC}(A_{\epsilon, \delta'p}^+) = \gamma_\epsilon^+(\delta'p)$, where

$$\begin{aligned}\gamma_\epsilon^+(\delta'p) &= p\delta'(1+\epsilon)\log(1+\epsilon) + (1-p\delta'(1+\epsilon))\log\left(\frac{1-p\delta'(1+\epsilon)}{1-p\delta'}\right) \\ &= p\delta'[(1+\epsilon)\log(1+\epsilon) - \epsilon] + o(p^2).\end{aligned}$$

Thus, as $p \rightarrow 0$, the rate decays linearly with p and as a consequence the sample size needed for a given precision increases proportionally to $1/p$.

Example 4.9 (Rare-event analysis for the light-tailed random walk). In the random walk example, Example 4.7, the performance of the importance sampling algorithm, with $\theta = \theta_a$, is captured by $\gamma_\epsilon^+(e^{m(\theta_a a - \kappa(\theta_a))}\delta)$. Taking $\delta = \delta'p_m$ where $p_m = \mathbb{P}(S_m \geq ma)$ it follows that

$$\begin{aligned}\gamma_\epsilon^+(e^{m(\theta_a a - \kappa(\theta_a))}\delta'p_m) &= e^{m(\theta_a a - \kappa(\theta_a))}p_m\delta'[(1+\epsilon)\log(1+\epsilon) - \epsilon] \\ &\quad + o(e^{2m(\theta_a a - \kappa(\theta_a))}p_m^2).\end{aligned}$$

The reduction in computational cost for the importance sampler vs. the standard Monte Carlo algorithm is then bounded from above by

$$c \frac{\gamma_\epsilon^+(\delta'p_m)}{\gamma_\epsilon^+(e^{m(\theta_a a - \kappa(\theta_a))}\delta'p_m)} \sim ce^{-m(\theta_a a - \kappa(\theta_a))}, \quad \text{as } m \rightarrow \infty.$$

The conclusion is that the reduction in computational cost is exponential in m .

We end this section by demonstrating the performance of the the so-called zero-variance change of measure [1, p. 127]. This choice of sampling distribution is optimal for estimating a probability in the sense that the variance of the estimator is zero and is often used as a reference point for designing efficient sampling distributions.

Example 4.10 (Zero-variance change of measure). Consider the probability $p = F(A)$, for some $A \subset \mathcal{X}$ and distribution F and take the importance function to be $I\{x \in A\}$. The zero-variance sampling distribution is the distribution \tilde{F} given by

$$\frac{d\tilde{F}}{dF}(x) = \frac{I\{x \in A\}}{p}.$$

Let $\delta' \in (0, 1)$ and $\delta = \delta'p$. The likelihood ratio is constant over the entire set A and $\tilde{F}(C) = F(C)p^{-1}$ for any $C \subset A$. That is, $\tilde{F}(C) = F(C | A)$, the conditional probability under F of C given A . It follows that

$$J_+(C) = \inf\{\mathcal{H}(G | \tilde{F}) : \frac{dG}{d\tilde{F}}(x) \geq 1 + \epsilon, x \in C\} = \gamma_\epsilon^+(\tilde{F}(C)) = \gamma_\epsilon^+(F(C)/p),$$

and the rate

$$I(A_{\epsilon, \delta'p}^+) = \gamma_\epsilon^+(\delta'),$$

which is independent of p .

For $\delta' = 1$, the case of having a relative error ϵ on the estimate of p , the rate corresponding to the zero-variance change of measure is $+\infty$. This follows from the fact that no probability measure G which is absolutely continuous with respect to \tilde{F} on A , and thus must satisfy $G(A) = 1$, can give rise to such an error. In this particular case the rate $(+\infty)$ can be obtained without using large deviations results.

4.3. Performance analysis for computing quantiles. Let the underlying space \mathcal{X} be the real line and consider computing a quantile of a distribution F on \mathcal{R} . For $\alpha \in (0, 1)$ the α -quantile of a finite measure ν on \mathcal{R} is defined by

$$\Phi_\alpha(\nu) = \inf\{x : \nu(x, \infty) \leq \alpha\}.$$

For $\epsilon > 0$, consider the set

$$A_\epsilon^+ = \{\nu \in \mathcal{M} : \Phi_\alpha(\nu) \geq (1 + \epsilon)\Phi_\alpha(F)\}.$$

Let \tilde{F} be the sampling distribution of an importance sampling algorithm for computing $\Phi_\alpha(F)$. Suppose that the importance function is an indicator $I\{x \in (a, \infty)\}$ where $a \leq \Phi_\alpha(F)$ and let I be the rate function in Theorem 3.1. The performance of the importance sampling algorithm can be quantified by $I(A_\epsilon^+)$. Of course one may also consider A_ϵ^- defined in the obvious way, but for this illustration we work exclusively with A_ϵ^+ .

Let us compute the rate $I(A_\epsilon^+)$. First note that, with $q_{\alpha, \epsilon} = (1 + \epsilon)\Phi_\alpha(F)$,

$$A_\epsilon^+ = \{\nu \in \mathcal{M} : \Phi_\alpha(\nu) \geq q_{\alpha, \epsilon}\} = \{\nu \in \mathcal{M} : \nu(q_{\alpha, \epsilon}, \infty) \geq \alpha\}$$

and the rate is therefore given by

$$\begin{aligned} I(A_\epsilon^+) &= \inf\{\mathcal{H}(G \mid \tilde{F}) : \nu = \Psi(G), \nu(q_{\alpha, \epsilon}, \infty) \geq \alpha\} \\ &= \inf\left\{\mathcal{H}(G \mid \tilde{F}) : \int_{\mathcal{R}} I\{x > q_{\alpha, \epsilon}\} w(x) G(dx) \geq \alpha\right\}. \end{aligned}$$

The infimum is attained at G_α^* given by

$$\frac{dG_\alpha^*}{d\tilde{F}}(x) = \frac{e^{\lambda k(x)}}{M(\lambda)},$$

where $k(x) = I\{x > q_{\alpha, \epsilon}\} w(x)$, $M(\lambda) = \int e^{\lambda k(x)} \tilde{F}(dx)$ and λ is given as the solution to

$$\alpha = \frac{\partial_\lambda M(\lambda)}{M(\lambda)}. \quad (4.3)$$

To see that the infimum is indeed attained at G_α^* , note that by the variational formula for relative entropy [4, Proposition 4.5.1], for all $\lambda \geq 0$ and $G \in \mathcal{M}_1$ such that $\int k(x) G(dx) \geq \alpha$,

$$\mathcal{H}(G \mid \tilde{F}) \geq \lambda \alpha - \log M(\lambda),$$

and the inequality is satisfied with equality for G_α^* . We have just proved the following.

Proposition 4.11. $I(A_\epsilon^+) = \lambda \alpha - \log M(\lambda)$ where λ is determined by (4.3).

Example 4.12. For a standard Monte Carlo algorithm the rate $I(A_\epsilon^+)$ can be explicitly computed. Indeed, in this case $k(x) = I\{x \geq q_{\alpha, \epsilon}\}$ and

$$M(\lambda) = e^\lambda p_{\alpha, \epsilon} + 1 - p_{\alpha, \epsilon},$$

with $p_{\alpha, \epsilon} = F(q_{\alpha, \epsilon}, \infty) \leq \alpha$. The equation for λ becomes

$$\alpha = \frac{\partial_\lambda M(\lambda)}{M(\lambda)} = \frac{e^\lambda p_{\alpha, \epsilon}}{e^\lambda p_{\alpha, \epsilon} + 1 - p_{\alpha, \epsilon}},$$

which leads to

$$\lambda = \log \left(\frac{\alpha(1 - p_{\alpha,\epsilon})}{(1 - \alpha)p_{\alpha,\epsilon}} \right)$$

and finally

$$I(A_\epsilon^+) = \lambda\alpha - \log M(\lambda) = \alpha \log \left(\frac{\alpha}{p_{\alpha,\epsilon}} \right) + (1 - \alpha) \log \left(\frac{1 - \alpha}{1 - p_{\alpha,\epsilon}} \right) = \mathcal{H}(\alpha \mid p_{\alpha,\epsilon}),$$

where $\mathcal{H}(\alpha \mid p_{\alpha,\epsilon})$ refers to the relative entropy between two Bernoulli distributions with parameters α and $p_{\alpha,\epsilon}$, respectively.

For a general importance sampling algorithm the expression for $I(A_\epsilon^+)$ in Proposition 4.11 has to be worked out on a case-by-case basis.

5. PROOF OF THEOREM 3.1

In this section the proof of the Laplace principle for the weighted empirical measures of importance sampling is presented. The proof relies on the weak convergence approach developed by Dupuis and Ellis [4]. The three main steps of the proof are:

- (1) Derive a representation formula for

$$W^n = -\frac{1}{n} \log \tilde{\mathbb{E}}[e^{-nh(\tilde{\mathbf{F}}_n^{wf})}].$$

This is achieved by formulating a stochastic control problem that has a minimal cost function which is equal to W^n . In the setting considered here, the representation formula reads

$$W^n = \inf_{\{G_{n,j}\}} \tilde{\mathbb{E}} \left[\frac{1}{n} \sum_{j=0}^{n-1} \mathcal{H}(G_{n,j}(\cdot \mid \bar{\mathbf{F}}_{n,j}) \mid \tilde{F}) + h(\bar{\mathbf{F}}_n) \right], \quad (5.1)$$

where $\bar{\mathbf{F}}_{n,j}$ is the controlled process (empirical measure), $\bar{\mathbf{F}}_{n,j+1} = \bar{\mathbf{F}}_{n,j} + \frac{1}{n} \delta_{\bar{X}_{n,j}}$ and $\bar{\mathbf{F}}_n = \frac{1}{n} \sum_{j=0}^{n-1} \delta_{\bar{X}_{n,j}}$, obtained by sampling $\bar{X}_{n,j}$ from the distribution (control) $G_{n,j}(\cdot \mid \bar{\mathbf{F}}_{n,j})$.

- (2) The representation formula (5.1) is used to prove the Laplace principle lower bound,

$$\liminf_n \frac{1}{n} \log \tilde{\mathbb{E}}[e^{-nh(\tilde{\mathbf{F}}_n^{wf})}] \geq - \inf_{G \in \Delta \cap \Gamma} \{h(\Psi(G)) + \mathcal{H}(G \mid \tilde{F})\}.$$

- (3) The third and most involved step is to use the representation formula to prove the Laplace principle upper bound,

$$\limsup_n \frac{1}{n} \log \tilde{\mathbb{E}}[e^{-nh(\tilde{\mathbf{F}}_n^{wf})}] \leq - \inf_{G \in \Delta \cap \Gamma} \{h(\Psi(G)) + \mathcal{H}(G \mid \tilde{F})\}.$$

The steps (1)-(3) are precisely those taken in [4] for proving Sanov's theorem. The main difficulty is to prove the upper bound, in step (3). Our proof of Theorem 3.1 is, for the most part, a transcription of the proof in [4]. The main new difficulties are that the mapping Ψ is defined on a subset of \mathcal{M}_1 and may be unbounded. The first point is handled by making minor adaptations to the arguments in [4] and the unboundedness of Ψ is mainly treated in Lemma 5.9 below. In order to make the paper self-contained results and constructions that are very similar to [4] are included and the corresponding references provided. In many cases the notation is consistent with that of [4], to make comparisons easier.

5.1. Representation formula. In this section (5.1) is shown to be equal to

$$W^n = -\frac{1}{n} \log \tilde{\mathbb{E}}[e^{-nh(\tilde{\mathbf{F}}_n^{wf})}].$$

The quantity W^n is the one that appears in the Laplace principle (with the sign changed) and (5.1) is derived by considering a related stochastic control problem described below. The corresponding minimal cost function can be used as a representation of W^n . For a more thorough discussion of this control problem see [4, Section 2.3].

The difference from Sanov's theorem is that only a subset $\Gamma = \{G \in \mathcal{M}_1 : G(wf) < \infty\}$ of the space of probability measures is under consideration. The proof is therefore very similar to the standard case. Recall that $\tilde{\mathbf{F}}_n^{wf}(g) = \tilde{\mathbf{F}}_n(wfg)$ for each measurable g . In particular,

$$\tilde{\mathbb{E}}[\tilde{\mathbf{F}}_n^{wf}(I_{\mathcal{X}})] = \tilde{\mathbb{E}}[\tilde{\mathbf{F}}_n(wf)] = F(f) < \infty,$$

and it follows that $\tilde{\mathbf{F}}_n \in \Gamma$ with probability 1. Let $\Gamma_n = \Gamma$ and define, for $j = 0, \dots, n-1$, recursively the sets $\Gamma_j \subset \mathcal{M}_{j/n}$ by

$$\Gamma_j = \{G \in \mathcal{M}_{j/n} : \tilde{F}(\{y : G + \frac{1}{n}\delta_y \in \Gamma_{j+1}\}) = 1\}.$$

That is, if $G \in \Gamma_j$, then sampling Y from \tilde{F} implies that $G + n^{-1}\delta_Y \in \Gamma_{j+1}$ with probability 1.

Let \mathcal{M}_0 be the one point set containing only the null measure and introduce the measurable mapping $W^n : \bigcup_{k=0}^n \{k\} \times \mathcal{M}_{k/n} \rightarrow \overline{\mathcal{R}} = [-\infty, \infty]$ by

$$W^n(j, G) = \begin{cases} -\frac{1}{n} \log \tilde{\mathbb{E}}[e^{-nh(\Psi(\tilde{\mathbf{F}}_n))} \mid \tilde{\mathbf{F}}_{n,j} = G], & \text{for } G \in \Gamma_j, \\ \infty, & \text{for } G \in \Gamma_j^c, \end{cases} \quad (5.2)$$

for $j = 0, \dots, n-1$ and

$$W^n(n, G) = \overline{h}(G), \quad (5.3)$$

where

$$\overline{h}(G) = \begin{cases} h(\Psi(G)), & \text{for } G \in \Gamma, \\ \infty, & \text{for } G \in \Gamma^c. \end{cases} \quad (5.4)$$

In particular, we set $W^n = W^n(0, 0)$. Note that in (5.2) $G \in \mathcal{M}_{j/n}(\mathcal{X})$ is a subprobability measure, and the $\tilde{\mathbf{F}}_{n,j} = (1/n) \sum_{i=0}^{j-1} \delta_{\tilde{X}_i}$ are the empirical subprobability measures obtained by sampling the \tilde{X}_i 's from \tilde{F} . Since the $\tilde{\mathbf{F}}_{n,j}$ form a Markov chain, one can obtain the recursion formula

$$W^n(j, G) = -\frac{1}{n} \log \int e^{-nW^n(j+1, G + \frac{1}{n}\delta_x)} \tilde{F}(dx).$$

Since h is bounded and continuous the mapping W^n is measurable, bounded from below, and bounded from above on Γ . Together with the recursion formula above, Proposition 4.5.1 in [4] gives that $W^n(j, G)$ can be written as

$$W^n(j, G) = \inf_{\tilde{G} \in \Delta} \left\{ \frac{1}{n} \mathcal{H}(\tilde{G} \mid \tilde{F}) + \int W^n(j+1, G + \frac{1}{n}\delta_x) \tilde{G}(dx) \right\}, \quad (5.5)$$

for $j = 0, \dots, n-1$, where

$$\Delta = \{G \in \mathcal{M}_1 : \mathcal{H}(G \mid \tilde{F}) < \infty\}.$$

Moreover, the infimum is attained at $\tilde{G}_{n,j} \in \Delta$ defined by the Radon-Nikodym derivative

$$\frac{d\tilde{G}_{n,j}}{d\tilde{F}}(x) = \frac{e^{-W^n(j+1, G + \frac{1}{n}\delta_x)}}{\int e^{-W^n(j+1, G + \frac{1}{n}\delta_y)} \tilde{F}(dy)}. \quad (5.6)$$

To derive the representation formula for W^n , consider the following related stochastic control problem. For $n \in \mathbb{N}$ and $j \in \{0, 1, \dots, n\}$, let $G_{n,j}$ be a stochastic kernel on \mathcal{X} given $\mathcal{M}_{j/n}$. A controlled process $\{\bar{\mathbf{F}}_{n,j}\}$ is defined by $\bar{\mathbf{F}}_{n,0} = \bar{\mathbf{F}}_{n,0} = 0$ and

$$\bar{\mathbf{F}}_{n,j} = \frac{1}{n} \sum_{k=0}^{j-1} \delta_{\bar{X}_{n,k}}, \quad \bar{\mathbf{F}}_n = \bar{\mathbf{F}}_{n,n},$$

where the conditional distribution of $\bar{X}_{n,k}$ given $\bar{\mathbf{F}}_{n,0}, \bar{\mathbf{F}}_{n,1}, \dots, \bar{\mathbf{F}}_{n,k}$ is

$$\bar{\mathbb{P}}(\bar{X}_{n,k} \in dx \mid \bar{\mathbf{F}}_{n,0}, \bar{\mathbf{F}}_{n,1}, \dots, \bar{\mathbf{F}}_{n,k}) = G_{n,j}(dx \mid \bar{\mathbf{F}}_{n,k}).$$

All random variables and the corresponding (controlled) empirical subprobability measures are for all n defined on a common probability space $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}})$ which will be used throughout the paper. For these dynamics define the minimal cost functions, for $j \in \{0, 1, \dots, n\}$ and $G \in \mathcal{M}_{j/n}$,

$$\bar{W}^n(j, G) = \inf_{\{G_{n,j}\}} \bar{\mathbb{E}}\left[\frac{1}{n} \sum_{k=j}^{n-1} \mathcal{H}(G_{n,k}(\cdot \mid \bar{\mathbf{F}}_{n,k}) \mid \tilde{F}) + h(\bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,j} = G\right]. \quad (5.7)$$

For $j = 0$ and $G = 0$ we set

$$\bar{W}^n = \bar{W}^n(0, 0) = \inf_{\{G_{n,j}\}} \bar{\mathbb{E}}\left[\frac{1}{n} \sum_{k=0}^{n-1} \mathcal{H}(G_{n,k}(\cdot \mid \bar{\mathbf{F}}_{n,k}) \mid \tilde{F}) + h(\bar{\mathbf{F}}_n)\right]. \quad (5.8)$$

Proposition 5.1. *Let \bar{W}^n be given by (5.7) and W^n be the solution to (5.2) and (5.3). Then $W^n = \bar{W}^n$.*

Proof. We begin by proving that $\bar{W}^n(j, G) \geq W^n(j, G)$ using backwards induction on j . Fix a control sequence $\{G_{n,j}\}$ and let $\{\bar{\mathbf{F}}_{n,j}\}$ denote the associated controlled process. Let $\bar{\Gamma}_n = \Gamma$ and define recursively for $j = 0, 1, \dots, n-1$ the sets $\bar{\Gamma}_j$ associated with the control sequence $\{G_{n,j}\}$ by

$$\bar{\Gamma}_j = \{G \in \mathcal{M}_{j/n} : G_{n,j}(\{y : G + \frac{1}{n}\delta_y \in \Gamma_{j+1}\}) = 1\}.$$

The definition is such that if at time j the controlled process $\bar{\mathbf{F}}_{n,j}$ lies in $\bar{\Gamma}_j$, then by sampling from $G_{n,j+1}, \dots, G_{n,n-1}$ the controlled process $\bar{\mathbf{F}}_n$ will belong to Γ with probability 1.

Consider first the case $j = n$. Clearly,

$$\bar{W}^n(n, G) = \bar{h}(G) = W^n(n, G),$$

and the claim is trivial for this j .

Take $j = n - 1$. Suppose that $G \in \bar{\Gamma}_{n-1}$, so that $G_{n,n-1}(\{y : G + \frac{1}{n}\delta_y \in \bar{\Gamma}_n\}) \mid G) = 1$. Using (5.2), (5.3) and (5.5), we have

$$\begin{aligned} & \mathbb{E}\left[\frac{1}{n}\mathcal{H}(G_{n,n-1}(\cdot \mid \bar{\mathbf{F}}_{n,n-1}) \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,n-1} = G\right] \\ &= \mathbb{E}\left[\frac{1}{n}\mathcal{H}(G_{n,n-1}(\cdot \mid \bar{\mathbf{F}}_{n,n-1}) \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n) + \int W^n(n, \bar{\mathbf{F}}_n) dG_{n,n-1} \right. \\ &\quad \left. - \int W^n(n, \bar{\mathbf{F}}_n) dG_{n,n-1} \mid \bar{\mathbf{F}}_{n,n-1} = G\right] \\ &\geq \mathbb{E}[W^n(n-1, \bar{\mathbf{F}}_{n,n-1}) + \bar{h}(\bar{\mathbf{F}}_n) - \int W^n(n, \bar{\mathbf{F}}_n) dG_{n,n-1} \mid \bar{\mathbf{F}}_{n,n-1} = G] \\ &= W^n(n-1, G). \end{aligned}$$

If instead $G \in \bar{\Gamma}_{n-1}^c$, then $G_{n,n-1}(\{y : G + \frac{1}{n}\delta_y \in \bar{\Gamma}_n\}) \mid G) < 1$ and since \bar{h} is infinite on $\bar{\Gamma}_n^c$, this implies

$$\mathbb{E}\left[\frac{1}{n}\mathcal{H}(G_{n,n-1}(\cdot \mid \bar{\mathbf{F}}_{n,n-1}) \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,n-1} = G\right] = \infty.$$

This shows that

$$\mathbb{E}\left[\frac{1}{n}\mathcal{H}(G_{n,n-1}(\cdot \mid \bar{\mathbf{F}}_{n,n-1}) \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,n-1} = G\right] \geq W^n(n-1, G),$$

for any choice of G .

Proceeding similarly for $j = n-2, n-3, \dots, 0$ then shows that

$$\mathbb{E}\left[\frac{1}{n} \sum_{k=j}^{n-1} \mathcal{H}(G_{n,k}(\cdot \mid \bar{\mathbf{F}}_{n,k}) \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,j} = G\right] \geq W^n(j, G),$$

for all G and j . Taking infimum over all admissible control sequences $\{G_{n,j}\} \subset \Delta$ proves the inequality.

Next, the reverse inequality $\bar{W}^n(j, G) \leq W^n(j, G)$ is proved. For this, consider the control sequence $\{\tilde{G}_{n,j}\}$ defined by (5.6). For this sequence it holds that $\bar{\Gamma}_j = \Gamma_j$ for all j .

The case $j = n$ was handled above and thus we start by considering $j = n-1$. If $G \in \Gamma_{n-1}$, then by the definition of $\tilde{G}_{n,n-1}$ and (5.5),

$$\begin{aligned} & \mathbb{E}\left[\frac{1}{n}\mathcal{H}(\tilde{G}_{n,n-1}(\cdot \mid \bar{\mathbf{F}}_{n,n-1}) \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,n-1} = G\right] \\ &= \mathbb{E}\left[\frac{1}{n}\mathcal{H}(\tilde{G}_{n,n-1}(\cdot \mid \bar{\mathbf{F}}_{n,n-1}) \mid \tilde{F}) + W^n(n, \bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,n-1} = G\right] \\ &= \mathbb{E}\left[\frac{1}{n}\mathcal{H}(\tilde{G}_{n,n-1}(\cdot \mid \bar{\mathbf{F}}_{n,n-1}) \mid \tilde{F}) \mid \bar{\mathbf{F}}_{n,n-1} = G\right] + \mathbb{E}[W^n(n, \bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,n-1} = G] \\ &= \frac{1}{n}\mathcal{H}(\tilde{G}_{n,n-1}(\cdot \mid G) \mid \tilde{F}) + \int W^n(n, G + \frac{1}{n}\delta_y) \tilde{G}_{n,n-1}(dy \mid G) \\ &= W^n(n-1, G). \end{aligned}$$

If instead $G \in \Gamma_{n-1}^c$, then $\tilde{G}_{n,n-1}(\{y : G + \frac{1}{n}\delta_y \in \Gamma\} \mid G) < 1$ and since \bar{h} is infinite on Γ^c we have

$$\mathbb{E}\left[\frac{1}{n}\mathcal{H}(\tilde{G}_{n,n-1}(\cdot \mid \bar{\mathbf{F}}_{n,n-1}) \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,n-1} = G\right] = \infty = W^n(n-1, G).$$

This shows that

$$\mathbb{E}\left[\frac{1}{n}\mathcal{H}(\tilde{G}_{n,n-1}(\cdot \mid \bar{\mathbf{F}}_{n,n-1}) \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,n-1} = G\right] = W^n(n-1, G),$$

for all G . Proceeding similarly for $j = n-2, \dots, 0$ shows that

$$\mathbb{E}\left[\frac{1}{n}\sum_{k=j}^{n-1}\mathcal{H}(\tilde{G}_{n,k}(\cdot \mid \bar{\mathbf{F}}_{n,k}) \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,j} = G\right] = W^n(j, G),$$

for all j and G . Taking infimum over all admissible control sequences $\{G_{n,j}\}$ in Δ yields the desired inequality. This completes the proof. \square

5.2. Laplace principle lower bound. In this section the Laplace principle lower bound,

$$\liminf_n \frac{1}{n} \log \mathbb{E}[e^{-nh(\bar{\mathbf{F}}_n^{wf})}] \geq - \inf_{G \in \Delta \cap \Gamma} \{h(\Psi(G)) + \mathcal{H}(G \mid \tilde{F})\}, \quad (5.9)$$

is proved. With W^n as in (5.2), proving this bound is equivalent to proving the upper bound

$$\limsup_n W^n \leq \inf_{G \in \Delta \cap \Gamma} \{h(\Psi(G)) + \mathcal{H}(G \mid \tilde{F})\}. \quad (5.10)$$

To this end the representation formula for W^n derived in Proposition 5.1,

$$W^n(j, G) = \inf_{\{G_{n,j}\}} \mathbb{E}\left[\frac{1}{n}\sum_{k=j}^{n-1}\mathcal{H}(G_{n,k}(\cdot \mid \bar{\mathbf{F}}_{n,k}) \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n) \mid \bar{\mathbf{F}}_{n,j} = G\right],$$

is used. The following strong law of large numbers will play a role in proving the lower bound (5.9).

Proposition 5.2 (Strong law of large numbers under importance sampling). *Let f be non-negative, measurable and F -integrable. Let $\{X_j\}$ be independent and identically distributed with common distribution F . Let \mathbf{F}_n^f be the weighted measure in \mathcal{M} determined by*

$$\mathbf{F}_n^f(g) = \frac{1}{n} \sum_{j=0}^{n-1} f(X_j)g(X_j),$$

for each bounded measurable function g . Then, with probability 1,

$$\mathbf{F}_n^f \xrightarrow{\tau} F^f,$$

in \mathcal{M} .

A proof for the corresponding result for empirical measures (i.e., no weights) is found in [4, pp. 49-50] and the proof for the case of weighted measures is a direct analogue.

Fix a probability measure $G \in \Delta \cap \Gamma$. Define the control sequence $\{G_{n,j}\}$ by $G_{n,j} = G$ for each $j = 0, 1, \dots, n-1$, so that in every step, the control does not depend on the controlled process. Then all the $\bar{\mathbf{X}}_{n,j}$'s are independent and identically distributed with common distribution G and the associated controlled process $\bar{\mathbf{F}}_n$ belongs to Γ with probability 1. Using the representation formula, it follows that

$$W^n \leq \mathbb{E}\left[\frac{1}{n}\sum_{k=0}^{n-1}\mathcal{H}(G \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n)\right] = \mathcal{H}(G \mid \tilde{F}) + \mathbb{E}[h(\Psi(\bar{\mathbf{F}}_n))].$$

The function wf is non-negative, measurable and G -integrable. By the choice G and how the controlled process $\bar{\mathbf{F}}_n$ is obtained, Proposition 5.2 implies that with probability 1, $\Psi(\bar{\mathbf{F}}_n; \cdot) \xrightarrow{\tau} \Psi(G; \cdot)$. Moreover, h is bounded and continuous with respect to the τ -topology and thus $h(\Psi(\bar{\mathbf{F}}_n)) \rightarrow h(\Psi(G))$ with probability 1. By the dominated convergence theorem,

$$\bar{\mathbb{E}}[h(\Psi(\bar{\mathbf{F}}_n))] \rightarrow h(\Psi(G)).$$

From this we conclude that for any $G \in \Delta \cap \Gamma$

$$\limsup_n W^n \leq \mathcal{H}(G \mid \tilde{F}) + h(\Psi(G)).$$

Finally, taking infimum over G in $\Delta \cap \Gamma$ on the right hand side proves the upper bound (5.10), and thus the Laplace principle lower bound.

5.3. Laplace principle upper bound. Analogously to the lower bound the Laplace principle upper bound,

$$\limsup_n \frac{1}{n} \log \mathbb{E}[e^{-nh(\bar{\mathbf{F}}_n^{wf})}] \leq - \inf_{G \in \Delta \cap \Gamma} [h(\Psi(G)) + \mathcal{H}(G \mid \tilde{F})], \quad (5.11)$$

can be stated as a lower limit for the minimal cost W^n ,

$$\liminf_n W^n \geq \inf_{G \in \Delta \cap \Gamma} [h(\Psi(G)) + \mathcal{H}(G \mid \tilde{F})]. \quad (5.12)$$

It therefore suffices to prove this lower limit for W^n in order to obtain the Laplace principle upper bound. To prove (5.12) it is enough to show that every subsequence has a further subsequence that satisfies the lower limit. Therefore, we henceforth work with a fixed subsequence also denoted by W^n .

Since $\mathcal{H}(\cdot \mid \tilde{F})$ is a convex function [4, Proposition 1.4.3],

$$\frac{1}{n} \sum_{j=0}^{n-1} \mathcal{H}(G_{n,j}(\cdot \mid \bar{\mathbf{F}}_{n,j}) \mid \tilde{F}) \geq \mathcal{H}\left(\frac{1}{n} \sum_{j=0}^{n-1} G_{n,j}(\cdot \mid \bar{\mathbf{F}}_{n,j}) \mid \tilde{F}\right),$$

and together with the representation formula this implies, with

$$\bar{G}_n = (1/n) \sum_{k=0}^{n-1} G_{n,j}(\cdot \mid \bar{\mathbf{F}}_{n,j}),$$

$$W^n \geq \inf_{\{G_{n,j}\}} \bar{\mathbb{E}}[\mathcal{H}(\bar{G}_n \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n)].$$

Given $\epsilon > 0$, there exists a control sequence $\{G_{n,j}\}$ and associated \bar{G}_n such that

$$W^n + \epsilon \geq \bar{\mathbb{E}}[\mathcal{H}(\bar{G}_n \mid \tilde{F}) + \bar{h}(\bar{\mathbf{F}}_n)] \quad (5.13)$$

There is no restriction on assuming $G_{n,j} \in \Delta \cap \Gamma$ for each j and for the remainder of this section this assumption is made.

To show the Laplace principle upper bound it is now enough to prove the following result.

Proposition 5.3. *Every subsequence of $\{(\bar{G}_n, \bar{\mathbf{F}}_n)\}$ has a further subsequence, also denoted $\{(\bar{G}_n, \bar{\mathbf{F}}_n)\}$, such that $(\Psi(\bar{G}_n), \Psi(\bar{\mathbf{F}}_n)) \xrightarrow{\tau} (\Psi(\bar{\mathbf{F}}), \Psi(\bar{\mathbf{F}}))$ with probability 1 along this subsequence, and where $\bar{\mathbf{F}}$ belongs to $\Delta \cap \Gamma$ with probability 1.*

Assume for now that Proposition 5.3 holds. The proof of the Laplace principle upper bound then follows from Fatou's lemma and the lower semi-continuity of the relative entropy mapping $G \mapsto \mathcal{H}(G \mid \tilde{F})$:

$$\begin{aligned} \epsilon + \liminf_n W^n &\geq \liminf_n \mathbb{E}[\mathcal{H}(\overline{G}_n \mid \tilde{F}) + \overline{h}(\overline{\mathbf{F}}_n)] \\ &\geq \mathbb{E}[\liminf_n \mathcal{H}(\overline{G}_n \mid \tilde{F}) + \liminf_n h(\Psi(\overline{\mathbf{F}}_n))] \\ &\geq \mathbb{E}[\mathcal{H}(\overline{\mathbf{F}} \mid \tilde{F}) + \overline{h}(\overline{\mathbf{F}})] \\ &\geq \inf_{G \in \Delta \cap \Gamma} \left\{ \mathcal{H}(G \mid \tilde{F}) + h(\Psi(G)) \right\}, \end{aligned}$$

where in the last step we used that $\overline{\mathbf{F}}$ is in $\Delta \cap \Gamma$ with probability 1.

Proposition 5.3 is proved by a series of lemmas. The idea is to first work with \mathcal{M}_1 equipped with the weak topology and show that $\{(\overline{G}_n, \overline{\mathbf{F}}_n)\}$ is tight in this topology. By showing that $\mathbb{E}[\mathcal{H}(\overline{G}_n \mid \tilde{F})]$ is uniformly bounded (Lemma 5.4) the tightness of $\{(\overline{G}_n, \overline{\mathbf{F}}_n)\}$ is obtained by showing tightness of each of the marginals. Prohorov's theorem implies relative compactness and thus each subsequence has a subsubsequence converging to some random element $(\overline{G}, \overline{\mathbf{F}})$. Lemma 5.6, which corresponds to [4, Lemma 2.5.1], concludes that $\overline{G} = \overline{\mathbf{F}}$ w.p. 1, thus establishing that w.p. 1 for each subsequence,

$$(\overline{G}_n, \overline{\mathbf{F}}_n) \xrightarrow{w} (\overline{\mathbf{F}}, \overline{\mathbf{F}}),$$

along some further subsequence. Next, in Lemma 5.7 it is established that the convergences of the marginals \overline{G}_n and $\overline{\mathbf{F}}_n$ to $\overline{\mathbf{F}}$ are still valid when \mathcal{M}_1 is equipped with the τ -topology. The main ingredient of the proof is an approximation argument introduced in [4, Lemma 9.3.3] and Lemma 5.7 below is a version of that result in the simpler setting where the underlying random variables are independent and identically distributed. Up to this point the proof is as in [4] with minor changes.

Once the convergence in \mathcal{M}_1 equipped with the τ -topology is established w.p. 1, it remains to show that it is preserved under the mapping Ψ . Lemma 5.8 proves that $\overline{\mathbf{F}}$ is in Γ w.p. 1 and thus that $\Psi(\overline{\mathbf{F}}; \cdot)$ is well-defined. The main additional difficulty is handled in Lemma 5.9 where a truncation argument is used to prove that $\Psi(\overline{G}_n; \cdot)$ converges to $\Psi(\overline{\mathbf{F}}; \cdot)$ and $\Psi(\overline{\mathbf{F}}_n; \cdot)$ converges to $\Psi(\overline{\mathbf{F}}; \cdot)$ in the τ -topology on \mathcal{M} .

Lemma 5.4. *For a sequence $\{\overline{G}_n\}$ of control sequences such that (5.13) holds and $G_{n,j} \in \Gamma \cap \Delta$ for each j and n , it holds that*

$$\sup_n \mathbb{E}[\mathcal{H}(\overline{G}_n \mid \tilde{F})] < \infty.$$

Proof. Since \overline{h} is bounded on Γ , it is possible to find a constant $M < \infty$ such that $\sup_{G \in \Gamma} |\overline{h}(G)| \leq M$. The choice of control sequence $\{G_{n,j}\}$ implies that $\overline{\mathbf{F}}_n \in \Gamma$ with probability 1 for all n . Therefore,

$$\begin{aligned} \sup_n \mathbb{E}[\mathcal{H}(\overline{G}_n \mid \tilde{F})] &= \sup_n \mathbb{E}[\mathcal{H}(\overline{G}_n \mid \tilde{F}) - M] + M \\ &\leq \sup_n \mathbb{E}[\mathcal{H}(\overline{G}_n \mid \tilde{F}) + \overline{h}(\overline{\mathbf{F}}_n)] + M \\ &\leq \sup_n (W^n + \epsilon) + M \\ &= \sup_n W^n + \epsilon + M. \end{aligned}$$

By the proof of the Laplace principle lower bound we have $\limsup_n W^n < \infty$, and together with $\mathbb{E}[\mathcal{H}(\overline{G}_n | \tilde{F})] < \infty$ for each n this implies that $\sup_n \mathbb{E}[\mathcal{H}(\overline{G}_n | \tilde{F})] < \infty$. \square

The uniform boundedness just established is used to prove the tightness of the sequence of admissible control measures.

Lemma 5.5. *Under condition (i) of Theorem 3.1, the sequence*

$$\left\{ \frac{1}{n} \sum_{i=0}^{n-1} G_{n,i}(\cdot | \overline{\mathbf{F}}_n) \times \overline{\mathbf{F}}_n \right\} = \{(\overline{G}_n \times \overline{\mathbf{F}}_n)\},$$

of admissible control measures in $\mathcal{M}_1 \times \mathcal{M}_1$ is tight in the weak topology.

Lemma 5.5 is a special case of Proposition 8.2.5 in [4] that establishes the result in the more general context of Markov chains. The proof is therefore omitted.

Having established the tightness of the distributions of $\{(\overline{G}_n, \overline{\mathbf{F}}_n)\}$, the next step is to extend this to almost sure convergence of subsequences, in the weak topology, and to study the limit.

Lemma 5.6. *Given any subsequence of $\{(\overline{G}_n, \overline{\mathbf{F}}_n)\}$, there exists a further subsequence that converges in distribution to some random variable $(\overline{G}, \overline{\mathbf{F}})$, where $\overline{G} = \overline{\mathbf{F}}$ a.s.*

The result in Lemma 5.6 is practically identical to part (b) of Lemma 2.5.1 in [4]. The only difference is that we must now appeal to the tightness proved in Lemma 5.5, whereas in [4] the underlying space is assumed to be compact. We omit the proof.

The next step is to show that this weak convergence of subsubsequences actually implies convergence of \overline{G}_n and $\overline{\mathbf{F}}_n$ in the τ -topology. The result is a version of Lemma 9.3.3 in [4] adapted to the case of independent and identically distributed random variables. Recall that we are already working with a specific subsubsequence (indexed by n) and on a probability space where the convergences $\overline{G}_n \xrightarrow{w} \overline{\mathbf{F}}$ and $\overline{\mathbf{F}}_n \xrightarrow{w} \overline{\mathbf{F}}$ both occur with probability 1. Henceforth, \mathcal{M}_1 will be equipped with the τ -topology.

Lemma 5.7. *Under the conditions of Lemmas 5.4-5.6, there exists some subsequence of $n \in \mathbb{N}$ such that $\overline{G}_n \xrightarrow{\tau} \overline{\mathbf{F}}$ and $\overline{\mathbf{F}}_n \xrightarrow{\tau} \overline{\mathbf{F}}$ w.p. 1 along this subsequence.*

An important ingredient in the proof is the inequality

$$ab \leq e^{a\sigma} + \frac{1}{\sigma}(b \log(b) - b + 1), \quad (5.14)$$

see the proof of [4, Lemma 9.3.3]. This inequality will also appear in what follows.

With the almost sure convergence in the τ -topology established the final results needed to prove Proposition 5.3 are obtained in Lemmas 5.8 and 5.9. The first result is that with probability 1 the limit measure $\overline{\mathbf{F}}$ is in the desired region of \mathcal{M}_1 .

Lemma 5.8. *Under the assumption $\int e^{\alpha w f} d\tilde{F} < \infty$ for any $\alpha > 0$, it holds that*

$$\sup_n \mathbb{E}[\overline{G}_n(wf)] < \infty, \text{ and } \sup_n \mathbb{E}[\overline{\mathbf{F}}_n(wf)] < \infty.$$

It follows that $\overline{\mathbf{F}} \in \Gamma$ w.p. 1.

Proof. Lemma 5.4 shows that $\sup_n \mathbb{E}[\mathcal{H}(\overline{G}_n \mid \tilde{F})] < \infty$. Hence, each \overline{G}_n has almost surely a well-defined Radon-Nikodym derivative w_n with respect to \tilde{F} and by definition

$$\mathcal{H}(\overline{G}_n \mid \tilde{F}) = \int w_n \log(w_n) d\tilde{F}.$$

Since w , f and w_n are all non-negative functions, the inequality (5.14) with $a = wf$, $b = w_n$ and $\sigma = 1$ gives

$$\begin{aligned} \overline{G}_n(wf) &= \int wf d\overline{G}_n = \int_{\mathcal{X}} wf w_n d\tilde{F} \\ &\leq \int e^{wf} d\tilde{F} + \int (w_n \log(w_n) - w_n - 1) d\tilde{F} \\ &= \int e^{wf} d\tilde{F} + \int \log(w_n) d\overline{G}_n \\ &= \int e^{wf} d\tilde{F} + \mathcal{H}(\overline{G}_n \mid \tilde{F}). \end{aligned}$$

Thus,

$$\mathbb{E}[\overline{G}_n(wf)] \leq \mathbb{E}[\mathcal{H}(\overline{G}_n \mid \tilde{F})] + \int e^{wf} d\tilde{F},$$

and

$$\sup_n \mathbb{E}[\overline{G}_n(wf)] \leq \sup_n \mathbb{E}[\mathcal{H}(\overline{G}_n \mid \tilde{F})] + \int e^{wf} d\tilde{F}.$$

By the assumption and Lemma 5.4 it holds that

$$\sup_n \mathbb{E}[\overline{G}_n(wf)] < \infty.$$

Lemma 5.7 proves that $\overline{G}_n \xrightarrow{\tau} \overline{\mathbf{F}}$ with probability 1. For $m \in \mathbb{N}$, $wf \wedge m$ is a bounded, measurable function and the τ -convergence implies that

$$\infty > \limsup_n \mathbb{E}[\overline{G}_n(wf)] \geq \limsup_n \mathbb{E}[\overline{G}_n(wf \wedge m)] = \mathbb{E}[\overline{\mathbf{F}}(wf \wedge m)].$$

Taking \limsup as $m \uparrow \infty$ together with a repeated use of Fatou's lemma gives $\mathbb{E}[\overline{\mathbf{F}}(wf)] < \infty$. This is a non-negative random variable and it follows that $\overline{\mathbf{F}} \in \Gamma$ a.s.

That also $\sup_n \mathbb{E}[\bar{\mathbf{F}}_n(wf)] < \infty$ is proved by a repeated conditioning on the controlled process.

$$\begin{aligned}
\mathbb{E}[\bar{\mathbf{F}}_n(wf)] &= \mathbb{E}\left[\frac{1}{n} \sum_{j=0}^{n-1} w(\bar{X}_{n,j}) f(\bar{X}_{n,j})\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\frac{1}{n} w(\bar{X}_{n,n-1}) f(\bar{X}_{n,n-1}) + \frac{1}{n} \sum_{j=0}^{n-2} w(\bar{X}_{n,j}) f(\bar{X}_{n,j}) \mid \bar{\mathbf{F}}_{n,n-1}\right]\right] \\
&= \mathbb{E}\left[\frac{1}{n} \int w(x) f(x) G_{n,n-1}(dx \mid \bar{\mathbf{F}}_{n,n-1}) + \bar{\mathbf{F}}_{n,n-1}\right] \\
&= \mathbb{E}\left[\frac{1}{n} G_{n,n-1}(wf \mid \bar{\mathbf{F}}_{n,n-1})\right. \\
&\quad \left.+ \mathbb{E}\left[\frac{1}{n} w(\bar{X}_{n,n-2}) f(\bar{X}_{n,n-2}) + \frac{1}{n} \sum_{j=0}^{n-3} w(\bar{X}_{n,j}) f(\bar{X}_{n,j}) \mid \bar{\mathbf{F}}_{n,n-2}\right]\right] \\
&= \mathbb{E}\left[\frac{1}{n} G_{n,n-1}(wf \mid \bar{\mathbf{F}}_{n,n-1}) + \frac{1}{n} G_{n,n-2}(wf \mid \bar{\mathbf{F}}_{n,n-2}) + \bar{\mathbf{F}}_{n,n-2}\right].
\end{aligned}$$

Proceeding like this one obtains

$$\begin{aligned}
\mathbb{E}[\bar{\mathbf{F}}_n(wf)] &= \mathbb{E}\left[\frac{1}{n} G_{n,0}(wf \mid \bar{\mathbf{F}}_{n,0}) + \dots + \frac{1}{n} G_{n,n-1}(wf \mid \bar{\mathbf{F}}_{n,n-1})\right] \\
&= \mathbb{E}\left[\frac{1}{n} \sum_{j=0}^{n-1} \int w(x) f(x) G_{n,j}(dx \mid \bar{\mathbf{F}}_{n,j})\right] \\
&= \mathbb{E}[\bar{G}_n(wf)].
\end{aligned}$$

This completes the proof. \square

Next it is shown that the almost sure convergence in the τ -topology on \mathcal{M}_1 implies almost sure convergence in the τ -topology on \mathcal{M} for the corresponding mapped measures $\Psi(\bar{G}_n; \cdot)$ and $\Psi(\bar{\mathbf{F}}_n; \cdot)$.

Lemma 5.9. *Along the subsequence for which the convergence in the τ -topology holds, the convergences*

$$\Psi(\bar{G}_n; \cdot) \xrightarrow{\tau} \Psi(\bar{\mathbf{F}}; \cdot), \text{ and } \Psi(\bar{\mathbf{F}}_n; \cdot) \xrightarrow{\tau} \Psi(\bar{\mathbf{F}}; \cdot),$$

in \mathcal{M} hold w.p. 1.

Proof. Define for $G \in \mathcal{M}_1$ and $m \in \mathbb{N}$ a truncated version $\Psi_m(G; \cdot)$ of the mapping Ψ as the finite measure given by

$$\Psi_m(G; g) = \int (wf \wedge m) g dG,$$

for each bounded, measurable function g . The function $wf \wedge m$ is bounded and measurable and thus Ψ_m is continuous with respect to the τ -topology on \mathcal{M}_1 . Therefore,

$$\Psi_m(\bar{G}_n; g) = \int (wf \wedge m) g d\bar{G}_n \rightarrow \int (wf \wedge m) g d\bar{\mathbf{F}} = \Psi_m(\bar{\mathbf{F}}; g),$$

as $n \rightarrow \infty$ along the particular subsequence. That is, for any bounded, measurable function g it holds w.p. 1 that

$$\lim_{n \rightarrow \infty} \Psi_m(\overline{G}_n; g) = \Psi_m(\overline{\mathbf{F}}; g).$$

Moreover, by Lemma 5.8 the function wf is a.s. integrable with respect to $\overline{\mathbf{F}}$. Since $(wf \wedge m)g \rightarrow wf g$ as $m \rightarrow \infty$, the dominated convergence theorem implies that with probability 1,

$$\lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} \Psi_m(\overline{G}_n; g) = \Psi(\overline{\mathbf{F}}; g),$$

for every bounded, measurable function g . Therefore, the desired convergence in \mathcal{M} with the τ -topology will follow if the order of the limit operators can be interchanged, which holds if

$$\lim_{m \rightarrow \infty} \sup_n \left| \int g w f d\overline{G}_n - \int g (w f \wedge m) d\overline{G}_n \right| = 0. \quad (5.15)$$

In proving Lemma 5.7, an application of the Skorohod representation theorem results in the sequence $\{\mathcal{H}(\overline{G}_n \mid \tilde{F})\}$ being bounded a.s. and (5.15) then follows by an application of the inequality (5.14); the argument goes precisely as in part (c) of Lemma 1.4.3 in [4]. This proves the part concerning the admissible control sequence.

To show convergence of $\Psi(\overline{\mathbf{F}}_n; \cdot)$ an argument similar to that of Lemma 5.8 is used. This line of reasoning is found in Lemmas 8.2.7. and 9.3.3. in [4]. The aim is to show that for each bounded measurable function g , $\Psi(\mathbf{F}_n; g) \rightarrow \Psi(\overline{\mathbf{F}}; g)$ in probability in a way such that for indicator functions $g = I_A$ we can appeal to the first Borel-Cantelli lemma to get almost sure convergence of the entire measure along some subsequence. For this, take any $\epsilon > 0$ and consider

$$\begin{aligned} & \mathbb{P}\left(\left|\int w f g d\overline{\mathbf{F}}_n - \int w f g d\overline{G}_n\right| \geq 3\epsilon\right) \\ & \leq \mathbb{P}\left(\left|\int (w f \wedge m) g d\overline{\mathbf{F}}_n - \int w f g d\overline{\mathbf{F}}_n\right| \geq \epsilon\right) \\ & \quad + \mathbb{P}\left(\left|\int (w f \wedge m) g d\overline{\mathbf{F}}_n - \int (w f \wedge m) g d\overline{G}_n\right| \geq \epsilon\right) \\ & \quad + \mathbb{P}\left(\left|\int (w f \wedge m) g d\overline{G}_n - \int w f g d\overline{G}_n\right| \geq \epsilon\right), \end{aligned}$$

which holds for any $m \geq 0$. For $n \in \mathbb{N}$ and $j \in \{0, 1, \dots, n-1\}$, define the σ -algebra generated by the controlled process up to time j ,

$$\overline{\mathcal{F}}_{n,j} = \sigma(\overline{\mathbf{F}}_{n,0}, \overline{\mathbf{F}}_{n,1}, \dots, \overline{\mathbf{F}}_{n,j}).$$

Recall that

$$\overline{\mathbb{P}}(\overline{X}_{n,j} \in dy \mid \overline{\mathbf{F}}_{n,0}, \overline{\mathbf{F}}_{n,1}, \dots, \overline{\mathbf{F}}_{n,j}) = G_{n,j}(dy \mid \overline{\mathbf{F}}_{n,j}),$$

that is $G_{n,j}(\cdot \mid \overline{\mathbf{F}}_{n,j})$ is a regular conditional distribution for $\overline{X}_{n,j}$ given $\overline{\mathcal{F}}_{n,j}$. Now condition on the $\overline{\mathcal{F}}_{n,j}$'s to relate expectation of integrals with respect to $\overline{\mathbf{F}}_{n,j}$ to

integrals with respect to the control measures $G_{n,j}$.

$$\begin{aligned}
& \mathbb{P}\left(\left|\int (wf \wedge m)gd\bar{\mathbf{F}}_n - \int wfgd\bar{\mathbf{F}}_n\right| \geq \epsilon\right) \\
& \leq \frac{1}{\epsilon}\mathbb{E}\left[\left|\int (wf \wedge m)gd\bar{\mathbf{F}}_n - \int wfgd\bar{\mathbf{F}}_n\right|\right] \\
& \leq \frac{\|g\|_\infty}{\epsilon}\mathbb{E}\left[\int (wf - wf \wedge m)d\bar{\mathbf{F}}_n\right] \\
& = \frac{\|g\|_\infty}{\epsilon}\mathbb{E}\left[\frac{1}{n}\sum_{j=0}^{n-1}(wf - wf \wedge m)(\bar{X}_{n,j})\right] \\
& = \frac{\|g\|_\infty}{\epsilon}\mathbb{E}\left[\frac{1}{n}\sum_{j=0}^{n-1}\mathbb{E}[(wf - wf \wedge m)(\bar{X}_{n,j}) \mid \bar{\mathcal{F}}_{n,j}]\right] \\
& = \frac{\|g\|_\infty}{\epsilon}\mathbb{E}\left[\frac{1}{n}\sum_{j=0}^{n-1}\int (wf - wf \wedge m)dG_{n,j}\right] \\
& = \frac{\|g\|_\infty}{\epsilon}\mathbb{E}\left[\int (wf - wf \wedge m)d\bar{G}_n\right].
\end{aligned}$$

Analogously, it holds that

$$\mathbb{P}\left(\left|\int (wf \wedge m)gd\bar{G}_n - \int wfgd\bar{G}_n\right| \geq \epsilon\right) \leq \frac{\|g\|_\infty}{\epsilon}\mathbb{E}\left[\int (wf - wf \wedge m)d\bar{G}_n\right].$$

Since $wf - wf \wedge m$ is non-negative and \bar{G}_n w.p. 1 has a Radon-Nikodym derivative w_n with respect to \bar{F} , we can apply the inequality (5.14). For any $\sigma \geq 1$,

$$\mathbb{E}\left[\int (wf - wf \wedge m)d\bar{G}_n\right] \leq \int e^{\sigma(wf - wf \wedge m)}d\bar{F} + \frac{1}{\sigma}\mathbb{E}[\mathcal{H}(\bar{G}_n \mid \bar{F})].$$

Finally, consider

$$\mathbb{P}\left(\left|\int (wf \wedge m)gd\bar{\mathbf{F}}_n - \int (wf \wedge m)gd\bar{G}_n\right| \geq \epsilon\right),$$

for which we use a similar argument as in the proof of [4, Lemma 8.2.7]. For any bounded measurable function $g : \mathcal{X} \rightarrow \mathcal{R}$,

$$\begin{aligned}
& \mathbb{E}[(wf \wedge m)g(\bar{X}_{n,j}) - \int (wf \wedge m)gdG_{n,j}(x \mid \bar{\mathbf{F}}_{n,j}) \mid \bar{\mathcal{F}}_{n,j}] \\
& = \mathbb{E}[(wf \wedge m)g(\bar{X}_{n,j}) \mid \bar{\mathcal{F}}_{n,j}] - G_{n,j}((wf \wedge m)g \mid \bar{\mathbf{F}}_{n,j}) \\
& = G_{n,j}((wf \wedge m)g \mid \bar{\mathbf{F}}_{n,j}) - G_{n,j}((wf \wedge m)g \mid \bar{\mathbf{F}}_{n,j}) = 0
\end{aligned}$$

\mathbb{P} -a.s. since $G_{n,j}(\cdot \mid \bar{\mathbf{F}}_{n,j})$ is a regular conditional distribution of $\bar{X}_{n,j}$. Hence,

$$\{(wf \wedge m)g(\bar{X}_{n,j}) - G_{n,j}((wf \wedge m)g \mid \bar{\mathbf{F}}_{n,j})\}_{j=0,1,\dots,n-1},$$

is a martingale difference sequence with respect to $\overline{\mathcal{F}}_{n,j}$. Moreover, for any $\epsilon > 0$,

$$\begin{aligned}
& \mathbb{P}\left(\left|\int (wf \wedge m)gd\overline{\mathbf{F}}_n - \int (wf \wedge m)gd\overline{\mathbf{G}}_n\right| \geq \epsilon\right) \\
& \leq \frac{1}{\epsilon^2} \mathbb{E}\left[\left|\int (wf \wedge m)gd\overline{\mathbf{F}}_n - \int (wf \wedge m)gd\overline{\mathbf{G}}_n\right|^2\right] \\
& = \frac{1}{\epsilon^2} \mathbb{E}\left[\frac{1}{n^2} \left(\sum_{j=0}^{n-1} ((wf \wedge m)g(\overline{X}_{n,j}) - G_{n,j}((wf \wedge m)g \mid \overline{\mathbf{F}}_{n,j}))\right)^2\right] \\
& = \frac{1}{\epsilon^2} \mathbb{E}\left[\frac{1}{n^2} \left(\sum_{j=0}^{n-1} ((wf \wedge m)g(\overline{X}_{n,j}) - G_{n,j}((wf \wedge m)g \mid \overline{\mathbf{F}}_{n,j}))^2\right.\right. \\
& \quad \left. + \sum_{i=1, j \neq i}^{n-1} ((wf \wedge m)g(\overline{X}_{n,i}) - G_{n,i}((wf \wedge m)g \mid \overline{\mathbf{F}}_{n,i})) \right. \\
& \quad \left. \times ((wf \wedge m)g(\overline{X}_{n,j}) - G_{n,j}((wf \wedge m)g \mid \overline{\mathbf{F}}_{n,j}))\right].
\end{aligned}$$

The second term vanishes when conditioning on the σ -algebras $\{\overline{\mathcal{F}}_{n,j}\}$. For the first term inside the expectation an upper bound on each of the summands is

$$\left((wf \wedge m)g(\overline{X}_{n,j}) - G_{n,j}((wf \wedge m)g \mid \overline{\mathbf{F}}_{n,j})\right)^2 \leq 4\|g\|_\infty^2 m^2.$$

It follows that

$$\begin{aligned}
& \frac{1}{\epsilon^2} \mathbb{E}\left[\frac{1}{n^2} \left(\sum_{j=0}^{n-1} ((wf \wedge m)g(\overline{X}_{n,j}) - G_{n,j}((wf \wedge m)g \mid \overline{\mathbf{F}}_{n,j}))\right)^2\right] \\
& \leq \frac{4\|g\|_\infty^2 m^2}{\epsilon^2 n}.
\end{aligned}$$

This all adds up to

$$\begin{aligned}
& \mathbb{P}\left(|\Psi(\overline{\mathbf{F}}_n; wfg) - \Psi(\overline{\mathbf{G}}_n; wfg)| \geq 3\epsilon\right) \\
& \leq \frac{\|g\|_\infty}{\epsilon} \left(\int e^{\sigma(wf - wf \wedge m)} d\tilde{F} + \frac{1}{\sigma} \mathbb{E}[\mathcal{H}(\overline{\mathbf{G}}_n \mid \tilde{F})]\right) + \frac{4\|g\|_\infty m^2}{\epsilon^2 n},
\end{aligned}$$

which goes to 0 if we send n, m and σ to $+\infty$ in that order. This holds due to the assumption $\int e^{\sigma wf} d\tilde{F} < \infty$, and Lemma 5.4.

To complete the proof we want to use the same argument as in Lemma 9.3.3 in [4] together with the convergence for $\Psi(\overline{\mathbf{G}}_n; \cdot)$. For this it is not enough to have the above convergence in probability but one needs to be able to find a subsequence for which the probabilities converge fast enough for an application of the first Borel-Cantelli lemma. Say that we want a subsequence n_k such that for $n \geq n_k$,

$$\mathbb{P}\left(|\Psi(\overline{\mathbf{F}}_n; wfg) - \Psi(\overline{\mathbf{G}}_n; wfg)| \geq 3\epsilon\right) \leq 2^{-k}.$$

Start by picking a σ_k such that

$$\frac{1}{\sigma_k} \sup_n \mathbb{E}[\mathcal{H}(\overline{\mathbf{G}}_n \mid \tilde{F})] \leq \frac{2^{-k}\epsilon}{3},$$

which is again possible due to Lemma 5.4. Having chosen this σ_k , pick m_k sufficiently large so that

$$\int e^{\sigma_k(wf - wf \wedge m_k)} d\tilde{F} \leq \frac{2^{-k}\epsilon}{3}.$$

This is possible due to the assumption $\tilde{F}(e^{\sigma wf}) < \infty$ for every $\sigma > 0$. Finally, pick n_k such that

$$\frac{m_k^2}{n_k} \leq \frac{2^{-k}\epsilon^2}{12}.$$

If g is an indicator function, the above will yield a sequence $\{n_k\}$ such that the probability of interest is smaller than 2^{-k} . The key to the argument used in Lemma 9.3.3 in [4] is that, since the underlying space \mathcal{X} is assumed complete and separable, the Borel σ -algebra on \mathcal{X} is generated by a countable collection of Borel sets. Together with the above this yields the desired result, namely that

$$\Psi(\overline{\mathbf{F}}_n; \cdot) \xrightarrow{\tau} \Psi(\mathbf{F}; \cdot),$$

in \mathcal{M} w.p. 1. We refer the reader to [4] for the details. \square

Lemma 5.4-5.9 completes the proof of Proposition 5.3.

We end this section by proving that the function I in (3.4) has sequentially compact level sets in the τ -topology.

Proposition 5.10. *The function $I : \mathcal{M} \rightarrow [0, \infty]$ in (3.4) has sequentially compact level sets on \mathcal{M} equipped with the τ -topology.*

Proof. Let

$$C(K) = \{\nu \in \mathcal{M} : I(\nu) \leq K\},$$

for $K < \infty$. Take any sequence $\{\nu_n\} \subset C(K)$. Since $I(\nu_j) \leq K$ for each j , there exists a sequence $\{G_n\} \subset \Delta \cap \Gamma$ such that $\Psi(G_j) = \nu_j$. Moreover, it must hold that $\mathcal{H}(G_j \mid \tilde{F}) \leq K + \epsilon$ for every $\epsilon > 0$ and each j . Hence,

$$\sup_j \mathcal{H}(G_j \mid \tilde{F}) < \infty.$$

The relative entropy has compact level sets in the τ -topology [4, Proposition 9.3.6]. Therefore, there exists some subsequence, also indexed by n , and some G_* such that

$$G_n \xrightarrow{\tau} G_*,$$

and the corresponding (finite) measures $\nu_n = \Psi(G_n)$, $\nu_* = \Psi(G_*)$ are in the set $C(K)$. Remains to prove that $\nu_n \xrightarrow{\tau} \nu_*$. To this end, we note that by the same arguments as in Lemma 5.8 it holds that

$$\sup_n G_n(wf) < \infty,$$

and

$$G_*(wf) < \infty.$$

The conditions used to prove Lemma 5.9 are therefore satisfied and it follows in the same way that for every bounded, measurable g ,

$$\nu_n(g) = \Psi(G_n; g) \rightarrow \Psi(G_*; g) = \nu_*(g)$$

as $n \rightarrow \infty$. Hence, $\nu_n \xrightarrow{\tau} \nu_*$ and the level set $C(K)$ is indeed compact in the τ -topology. \square

REFERENCES

- [1] S. Asmussen and P. W. Glynn. *Stochastic simulation: algorithms and analysis*. Stochastic Modelling and Applied Probability, 57. Springer, New York, 2007. MR2331321
- [2] P. Del Moral and J. Garnier. Genealogical particle analysis of rare events. *Ann. Appl. Probab.*, 15 (2005) no. 4, 2496–2534. MR2187302
- [3] A. Dembo and O. Zeitouni. *Large deviations techniques and applications*. 2nd ed. Applications of Mathematics (New York), 38. Springer-Verlag, New York, 1998. MR1619036
- [4] P. Dupuis and R. S. Ellis. *A weak convergence approach to the theory of large deviations*. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., New York, 1997. MR1431744
- [5] P. L’Ecuyer, V. Demers, and B. Tuffin. Splitting for rare event simulation. In *Proceedings of the 2006 Winter Simulation Conference*, 137–148. Winter Simulation Conference. IEEE Press, 2006.

(H. Hult) DEPARTMENT OF MATHEMATICS, KTH, 100 44 STOCKHOLM, SWEDEN
E-mail address: hult@kth.se

(P. Nyquist) DEPARTMENT OF MATHEMATICS, KTH, 100 44, STOCKHOLM, SWEDEN
E-mail address: pierren@kth.se